Randomized Matrix Decompositions using R

N. Benjamin Erichson
University of St Andrews

Sergey Voronin
Tufts University

Steven L. Brunton
University of Washington

J. Nathan Kutz
University of Washington

arXiv:1608.02148v2 [stat.CO] 3 Sep 2016

Abstract

The singular value decomposition (SVD) is among the most ubiquitous matrix factorizations. Specifically, it is a cornerstone algorithm for data analysis, dimensionality reduction and data compression. However, despite modern computer power, massive datasets pose a computational challenge for traditional SVD algorithms. We present the R package **rsvd**, which enables the fast computation of the SVD and related methods, facilitated by randomized algorithms. The **rsvd** package provides routines for computing the randomized singular value decomposition, randomized principal component analysis and randomized robust principal component analysis. Randomized algorithms provide an efficient computational framework for reducing the computational demands of traditional (deterministic) matrix factorizations. The key idea is to compute a compressed representation of the data using random sampling. This smaller (compressed) matrix captures the essential information that can then be used to obtain a low-rank matrix approximation. Several numerical examples demonstrate the usage of the **rsvd** package. The results show substantial accelerated computational times using the powerful concept of randomization.

Keywords: randomized singular value decomposition, randomized principal component analysis, robust principal component analysis, randomized subspace learning, R.

1. Introduction

Advances in sensor technology and data acquisition generate torrential streams of data. Massive datasets emerge across the social, physical, biological and ecological sciences, in the form of social networks, hyper spectral imagery, DNA microarrays, and animal movement data. This forces a shift from classical statistical data analysis concerned with a moderate size of observations and a few carefully selected and meaningful variables, towards the analysis of massive unstructured data (Donoho 2000).

Matrix factorizations are fundamental tools for data processing in statistical computing and machine learning. In particular, the singular value decomposition (SVD) is a cornerstone algorithm extensively used for data analysis, dimensionality reduction, and data compression. SVD is the workhorse algorithm behind linear regression, (robust) principal component analysis, discriminant analysis and canonical correlation analysis. However, the emergence of massive data poses a computational challenge for classical deterministic matrix algorithms. The

field of randomized numerical linear algebra provides a powerful framework to tackle this computational challenge (Drineas and Mahoney 2016). The key idea of randomized accelerated matrix algorithms is to compute a smaller representation of the original data matrix using randomization as a computational resource. Thus, traditional matrix decompositions can be obtained on a much smaller (compressed) matrix. This even enables decompositions of massive matrices for which traditional deterministic algorithms fail. Randomized algorithms are shown to be robust, while providing a substantial reduction of the computational demands (Mahoney 2011). This concept is most interesting because many high-dimensional data exhibit a low-dimensional structure, i.e., the intrinsic rank of the data matrix is much lower than the ambient dimension. Further, randomized algorithms allow to better exploit modern computational architectures, e.g., multithreading or parallelized and distributed computing.

Motivation and contributions. The computational costs of deterministic matrix algorithms for massive data matrices in R are tremendous. The singular value decomposition function `svd()` calls routines from the underlying LAPACK software package (Anderson, Bai, Bischof, Blackford, Dongarra, Du Croz, Greenbaum, Hammarling, McKenney, and Sorensen 1999). These routines are numerical stable and highly accurate (full double precision). However, in many applications, the full SVD is not necessary; only the truncated factorization is required. Using the full SVD to compute the truncated factorization is too expensive. The alternatives for constructing low-rank approximations of large-scale data are partial singular value decomposition algorithms. As far as we are aware, only the **svd** (Korobeynikov and Larsen 2016) and **Rspectra** (Qiu, Mei, Guennebaud, and Niesen 2016) packages provide competitive functions for computing the partial singular value decomposition in R. The former package provides a wrapper for the **PROPACK** SVD algorithm (Larsen 1998). The latter package is inspired by the software package **ARPACK** (Lehoucq, Sorensen, and Yang 1998) and provides a fast partial SVD and eigendecompositon algorithm. Both packages are in particular powerful for computing the singular value decomposition of large and sparse or structured matrices.

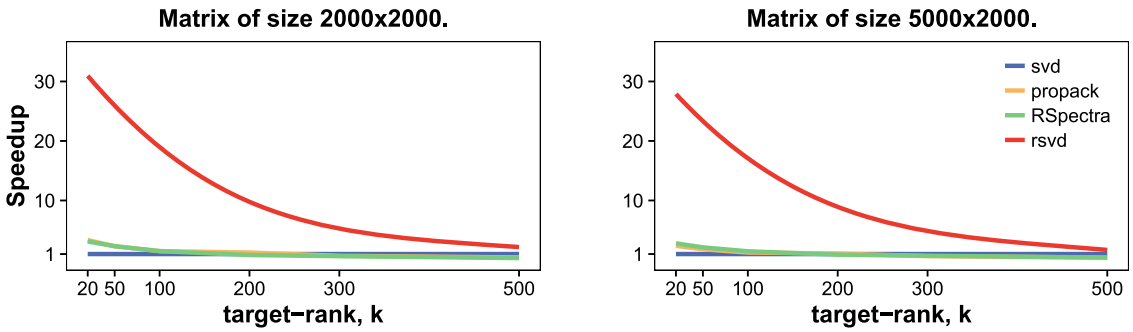

Figure 1: Computational performance of singular value decompositions in R.

Our package **rsvd** provides very efficient routines (accelerated via randomized sampling) for computing the partial singular value decomposition and for principal component analysis. Specifically, the randomized singular value decomposition function `rsvd()` gains a very significant speedup when one seeks a low-rank approximation with a small rank relative to the

ambient matrix dimension. Figure 1 shows the computational performance of different SVD routines for dense matrices in R, where k denotes the target-rank for the low-rank matrix approximation. Here the `svd()` function is compared to the `propack.svd()` of the **svd** package, the `svds()` of the **Rspectra** package, and to `rsvd()` function of the **rsvd** package. The randomized singular value decomposition outperforms, while attaining similar reconstruction errors as shown in Section 4. In Figure 1, the randomized algorithm is favorable for all target-ranks $k < \frac{min(m,n)}{4}$, where m and n denote the matrix dimensions. Specifically, the **rsvd** package[1] provides the following core functions:

- Randomized singular value decomposition: `rsvd()`.
- Randomized principal component analysis: `rpca()`.
- Randomized robust principal component analysis: `rrpca()`.

The interface of the `rsvd()` and `rpca()` functions is designed similar to the corresponding R base functions `svd()` and `prcomp()`. The latter function (based on the singular value decomposition) is the preferred routine for computing the principal component analysis in R. This is because, using the SVD is computational more efficient in general then using the eigenvalue decomposition, like the `princomp()` function does Venables and Ripley (2002). The `rpca()` can be used as a plug-in function for the `prcomp()` functions to benefit of the fast randomized algorithm, while providing the familiar summary statistics and plots. In addition, also fancy plot functions are provided to visualize the results of the principal component analysis. The `rrpca()` functions computes the robust principal component analysis using the augmented Lagrange multiplier method Lin, Chen, and Ma (2010). This function is designed so that the user can choose between either the randomized or deterministic SVD algorithm.

Organization. The remainder of this paper is organized as follows. Following a discussion below on notation, Section 2 briefly reviews the singular value and eigen decompositions and principal component analysis. Section 3 describes the fast randomized algorithm for computing the near optimal low-rank SVD in detail, followed by a discussion of different measurement matrices. This section also discusses the application of randomized methods to eigendecomposition and principal and robust principal component analysis. Section 4 presents motivating examples for using the **rsvd** package, including examples of image compression, eigenfaces, and foreground/background decomposition. Finally, the computational performance is evaluated on both dense and sparse matrices. Concluding remarks and a roadmap for further developments of the **rsvd** package are outlined in Section 5.

Notation. $\mathbf{A}$ denotes an $m \times n$ real valued matrix, and $\mathbf{A}^\top$ denotes the corresponding $n \times m$ transposed matrix. $\|\cdot\|$ denotes the spectral or operator norm, and $\|\cdot\|_F$ denotes the Frobenius norm, where $\|\mathbf{A}\|_F = \sqrt{\text{trace}(\mathbf{A}^\top \mathbf{A})}$. The relative reconstruction error is computed as $\|\mathbf{A} - \hat{\mathbf{A}}\|_F / \|\mathbf{A}\|_F$. We use $\mathsf{E}[\cdot]$ to denote the expected value and $\mathsf{VAR}[\cdot]$ to denote the variance of a random variable. $D = \text{diag}(a_1, \ldots, a_n)$ refers to an $n \times n$ diagonal matrix with $D_{(i,i)} = a_i$. The range (column) space of $\mathbf{A}$ is denoted as $R(\mathbf{A})$. By $\text{trunc}(\mathbf{U}, k)$ we refer to extracting the first or last k columns of $\mathbf{U}$.

[1] The project page is https://github.com/Benli11/rSVD.

2. Matrix decompositions and applications

Here we briefly describe the singular value and eigenvalue decomposition and their application to principal component and robust principal component analysis. For a comprehensive technical overview we refer to Golub and Van Loan (1996), Demmel (1997), and Watkins (2002).

Brief overview of the development of the singular value decomposition. While the origins of the SVD can be traced back to the late 19th century, the field of randomized matrix algorithms is relatively young. Figure 2 shows an incomplete time-line of some major developments of the singular value decomposition. Stewart (1993) gives an excellent historical review of the five mathematicians who developed the fundamentals of the SVD, namely Eugenio Beltrami (1835-1899), Camille Jordan (1838-1921), James Joseph Sylvester (1814-1897), Erhard Schmidt (1876-1959) and Hermann Weyl (1885-1955). The development and fundamentals of modern high-performance algorithms to compute the SVD is related to the seminal work of Golub and Kahan (1965) and Golub and Reinsch (1970).

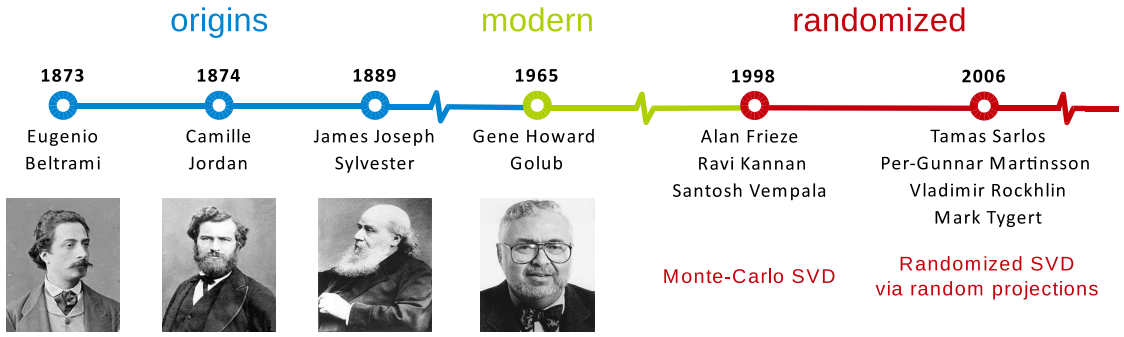

Figure 2: A timeline of major singular value decomposition developments.

Modern singular value decomposition algorithms are largely based on Krylov methods. These methods are accurate and in particular powerful for approximating structured or sparse large-scale matrices (Martinsson 2016). A modern and prominent state-of-the-art algorithm, based on the Lanczos algorithm, is the **PROPACK** SVD algorithm (Larsen 1998). As well as the partial SVD (`svds`) algorithm based on the **ARPACK** software package (Lehoucq *et al.* 1998), used in MATLAB or Python for sparse matrices.

Randomized matrix algorithms for computing low-rank matrix approximations have gained prominence over the past two decades. Frieze, Kannan, and Vempala (2004) introduced the 'Monte Carlo' SVD, a rigorous approach to efficiently compute the approximate low-rank SVD based on non-uniform row and column sampling. Sarlos (2006) and Martinsson, Rokhlin, and Tygert (2011) introduced a more robust approach based on random projections. Specifically, the properties of random vectors are exploited to efficiently build a subspace that captures the column space of a matrix. Woolfe, Liberty, Rokhlin, and Tygert (2008) further improved the computational performance by leveraging the properties of highly structured matrices which enable fast matrix multiplications. Eventually, the seminal work by Halko, Martinsson, and Tropp (2011b) unified and expanded the work on the randomized singular value decomposition (rSVD) and introduced state-of-the-art prototype algorithms to compute the near-optimal low-rank singular value decomposition.

2.1. Singular value and eigen decompositions.

The singular value decomposition (SVD) is among the most ubiquitous matrix factorizations of the computational era. The SVD is the workhorse algorithm for a wide variety of learning algorithms and data methods. In particular, the SVD provides a numerically stable matrix decomposition that can be used to obtain low-rank approximations, to compute the pseudo-inverses of non-square matrices, and to find the least-squares and minimum norm solutions of a linear model. Given an arbitrary real matrix[2] $\mathbf{A} \in \mathbb{R}^{m \times n}$, where $m \geq n$ without loss of generality, we seek a decomposition, such that

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top \tag{1}$$

as illustrated in Figure 3.

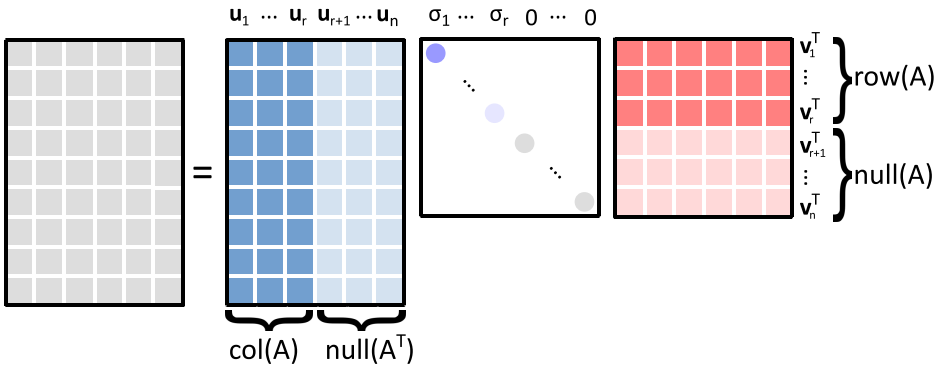

Figure 3: Schematic of the thin singular value decomposition.

The matrices $\mathbf{U} = [\mathbf{u}_1, ..., \mathbf{u}_m] \in \mathbb{R}^{m \times m}$ and $\mathbf{V} = [\mathbf{v}_1, ..., \mathbf{v}_n] \in \mathbb{R}^{n \times n}$ are orthonormal so that $\mathbf{U}^\top\mathbf{U} = \mathbf{I}$ and $\mathbf{V}^\top\mathbf{V} = \mathbf{I}$. The first r left singular vectors in $\mathbf{U}$ provide a basis for the range and the first r right singular vectors in $\mathbf{V}$ a basis for the domain of the matrix $\mathbf{A}$, whereby r denotes the rank of the data matrix. The rectangular diagonal matrix $\mathbf{S} \in \mathbb{R}^{m \times n}$ contains the corresponding non-negative singular values $\sigma_1 \geq ... \geq \sigma_r$, describing the spectrum of the data. If the rank r of the matrix $\mathbf{X}$ is smaller then the number of columns (i.e., $r < n$), then the last $m - r$ singular values $\{\sigma_i : i \geq r + 1\}$ are zero. The so called 'economic' or 'compact' SVD computes only the singular vectors corresponding to the non-zero singular values $\mathbf{U} = [\mathbf{u}_1, ..., \mathbf{u}_r] \in \mathbb{R}^{m \times r}$ and $\mathbf{V} = [\mathbf{v}_1, ..., \mathbf{v}_r] \in \mathbb{R}^{n \times r}$ respectively.

In many cases, the numerical rank $\bar{r}$ of the matrix is smaller than it's mathematical rank. That is, many of the last $\min(m, n) - \bar{r}$ singular values can be close to machine precision. Since the corresponding singular vectors are not in the span of the data, it is often desirable to compute only a reduced version of the SVD. The matrix can be well approximated by including only those eigenvectors which correspond to eigenvalues of a significant amplitude. This number k can be much smaller than $\min(m, n)$ depending on the value of $\bar{r}$. Choosing an optimal target rank k is highly dependent on the task, i.e. whether one is interested in a highly accurate reconstruction of the original data or in a very low dimensional representation of dominant

[2] Without loss of generality, the concept applies to complex matrices using the Hermitian transpose instead.

features in the data. The low-rank SVD of rank k takes the form:

$$\mathbf{A}_k = \mathbf{U}_k\mathbf{\Sigma}_k\mathbf{V}_k = [\mathbf{u}_1, ..., \mathbf{u}_k]\text{diag}(\sigma_1, \ldots, \sigma_k)[\mathbf{v}_1, ..., \mathbf{v}_k]^\top. \tag{2}$$

The Eckart-Young theorem (Eckart and Young 1936) states that the low-rank SVD provides the optimal rank-k reconstruction of a matrix in the least-square sense, both in the spectral and Frobenius norms:

$$\mathbf{A}_k = \mathbf{U}_k\mathbf{\Sigma}_k\mathbf{V}_k^\top := \underset{\mathbf{A}'_k}{\text{argmin}}\, \|\mathbf{A} - \mathbf{A}'_k\|. \tag{3}$$

Truncating small singular values in the the deterministic SVD gives an optimal approximation of the corresponding rank. The reconstruction error in both the spectral and Frobenius norms is given by:

$$\|\mathbf{A} - \mathbf{A}_k\| = \sigma_{k+1} \quad \text{and} \quad \|\mathbf{A} - \mathbf{A}_k\|_F = \sqrt{\sum_{j=k+1}^{\min(m,n)} \sigma_j^2}. \tag{4}$$

Closely related to the SVD is the eigendecomposition. When $\mathbf{M}$ is a square $n \times n$ matrix, we can write $\mathbf{M} = \mathbf{Q}\mathbf{\Sigma}\mathbf{Q}^{-1}$ where $\mathbf{Q}$ is the matrix of eigenvectors and $\mathbf{\Sigma}$ is the diagonal matrix of eigenvalues. Notice that for any matrix $\mathbf{A}$, with SVD given by $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, we have the eigendecompositions $\mathbf{A}^\top\mathbf{A} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^\top$ and $\mathbf{A}\mathbf{A}^\top = \mathbf{U}\mathbf{\Sigma}^2\mathbf{U}^\top$. These formulations are useful when $\mathbf{A}$ is a tall skinny or a short wide matrix. In the case where $\mathbf{\Sigma}$ contains no zero singular values, the SVD of $\mathbf{A}$ can be more efficiently computed via the eigendecomposition of a smaller $\mathbf{A}^\top\mathbf{A}$ or $\mathbf{A}\mathbf{A}^\top$ matrix. Notice, for instance that once $\mathbf{U}$ is computed, one can compute $\mathbf{V}$ using $\mathbf{A}^\top$ and the inverse of $\mathbf{\Sigma}$:

$$\mathbf{A}^\top\mathbf{U} = \mathbf{V}\mathbf{\Sigma} \Rightarrow \mathbf{V} = \mathbf{A}^\top\mathbf{U}\mathbf{\Sigma}^{-1}$$

Via a similar relation, one can compute $\mathbf{U}$ given $\mathbf{V}$ using $\mathbf{U} = \mathbf{A}\mathbf{V}\mathbf{\Sigma}^{-1}$.

2.2. Principal component analysis.

Originally formulated by Pearson (1901), principal component analysis (PCA)[3] still plays an important role in modern statistics due to its simple geometric interpretation. Specifically, it is widely used for feature extraction and visualization of big datasets comprising many interrelated variables. A classical statistical text on PCA is Jolliffe (2002), while more modern views and extensions are presented by Hastie, Tibshirani, and Friedman (2009), Murphy (2012) and Izenman (2008). Further, we want to point out the excellent review by Abdi and Williams (2010) and the recent seminal paper on linear dimensionality reduction by Cunningham and Ghahramani (2015).

The essential idea of PCA is to find a new set of uncorrelated variables that retain most of the information (total variation) present in the data. Figure 4 illustrates this for a two-dimensional example. Subplot (a) shows some fairly correlated data. The arrows indicate the principal directions of the data. It can be seen that the first arrow points into the direction which explains most of the variance, while the second arrow is orthogonal to the first one. Together, they span a new coordinate system so that the first axis accounts for most of the variance and the second for the remaining variance in the data. Subplot (b) shows the original data in

[3] Also commonly known as Hotelling transform, Karhunen Loève, or proper orthogonal decomposition (POD).

this new coordinate system, represented by a new set of uncorrelated variables the so called principal components. The values of these new variables are called principal component scores or coordinates. The histograms indicate that most information (variation) is now captured by just the first principal component (PC).

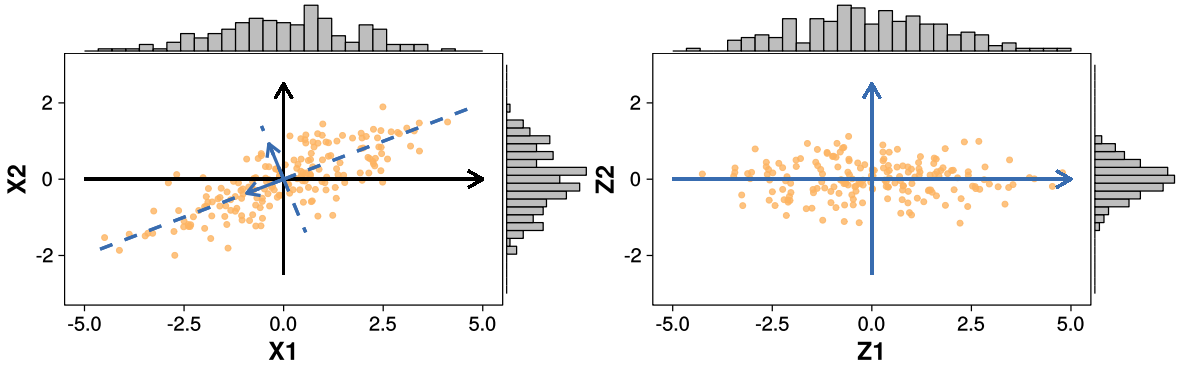

Figure 4: Illustration of PCA seeking to find a new set of uncorrelated variables. In (a) data and its two principal directions are shown. In (b) the new principal component variables are shown, indicating that the first component accounts for most of the variation in the data.

Given a data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ with n observations and p variables (column wise mean centered), the principal components can be expressed as a linear combination

$$\mathbf{z}_i = \mathbf{X}\mathbf{w}_i \tag{5}$$

where $\mathbf{z}_i \in \mathbb{R}^n$ denotes the the i'th principle component and the i'th principal direction is represented by the vector $\mathbf{w}_i \in \mathbb{R}^p$, where the elements of $\mathbf{w}_i = (w_1, ..., w_p)^\top$ are the corresponding principal component coefficients or weights.

In summary, we seek that the first principal component explains most of the total variation in the data. Subsequent PCs are required to be orthogonal to the previous components, and capturing the remaining variance in descending order. Mathematically, this problem can be formulated either as a least square minimization or as a variance maximization problem (Cunningham and Ghahramani 2015). However, both reduce to an eigenvalue problem. We follow the latter, and the more modern approach, i.e., maximizing the variance of the first principal component subject to the normalization constraint $\|\mathbf{w}\|_2^2 = 1$ as follows

$$\mathbf{w}_1 = \underset{\|\mathbf{w}\|_2^2=1}{\operatorname{argmax}} \mathsf{VAR}(\mathbf{X}\mathbf{w}) \propto \underset{\|\mathbf{w}\|_2^2=1}{\operatorname{argmax}} \|\mathbf{X}\mathbf{w}\|_2^2, \tag{6}$$

where $\|\cdot\|_2$ denotes the ℓ_2 norm. The last term can then be expanded as

$$\mathbf{w}_1 = \underset{\|\mathbf{w}\|_2^2=1}{\operatorname{argmax}} \mathbf{w}^\top (\mathbf{X}^\top \mathbf{X}) \mathbf{w}. \tag{7}$$

Considering that we constrained $\|\mathbf{w}\|_2^2 = \mathbf{w}^\top \mathbf{w} = 1$ to be a unit vector, Eq. (7) can be rewritten as

$$\mathbf{w}_1 = \underset{\mathbf{w}}{\operatorname{argmax}} \frac{\mathbf{w}^\top \mathbf{C} \mathbf{w}}{\mathbf{w}^\top \mathbf{w}}, \tag{8}$$

where $\mathbf{C}$ is a symmetric positive definite matrix, e.g., the covariance or correlation matrix. Specifically, the sample covariance matrix is defined as

$$\mathbf{C} = \frac{1}{n-1}\mathbf{X}^\top\mathbf{X}. \tag{9}$$

Finally, the method of Lagrange multipliers can be used to solve the problem

$$\mathcal{L}(\mathbf{w}_1, \lambda_1) = \max_{\mathbf{w}_1,\lambda_1}(\mathbf{w}_1^\top\mathbf{C}\mathbf{w}_1 - \lambda_1(\mathbf{w}_1^\top\mathbf{w}_1 - 1)), \tag{10}$$

which leads to the well known eigenvalue problem

$$\mathbf{C}\mathbf{w}_1 = \lambda_1\mathbf{w}_1. \tag{11}$$

Hence, the first principle direction for the mean centered matrix $\mathbf{X}$ is given by the dominant eigenvector $\mathbf{w}_1$. More generally, the principle directions of the covariance matrix $\mathbf{C}$ are the columns of the eigenvector matrix $\mathbf{W} \in \mathbb{R}^{p\times p}$ and the corresponding eigenvalues λ are the diagonal elements of $\mathbf{\Lambda} \in \mathbb{R}^{p\times p}$. Interestingly, the eigenvalues express exactly the amount of variation explained by the principal components.

$$\mathbf{C}\mathbf{W} = \mathbf{\Lambda}\mathbf{W}. \tag{12}$$

More compactly, we can compute the principal components $\mathbf{Z} \in \mathbb{R}^{n\times p}$ as

$$\mathbf{Z} = \mathbf{X}\mathbf{W}. \tag{13}$$

Hence, the matrix $\mathbf{W}$ can also be interpreted as a projection matrix that maps the original observations to the new coordinates in the eigenspace. Since the eigenvectors have unit norm the projection should be purely rotational without any scaling; thus, the matrix $\mathbf{W}$ is also denoted as a rotation matrix. However, we want to stress that the term 'loading' in general refers to the scaled eigenvectors

$$\mathbf{L} = \mathbf{W}\mathbf{\Sigma}^{0.5} \tag{14}$$

which, in some situations, provide a more insightful interpretation of the principal components.[4] The loading matrix $\mathbf{L} \in \mathbb{R}^{n\times p}$ has the properties: a) the squared column sums equal the eigenvalues, and b) the squared row sums equal the amount of a variable's variance.

In practice, we are often interested in a useful low-dimensional representation to reveal the coherent structure of the data. However, the number of principal components to retain is subtle and often domain specific. Many different heuristics (e.g. scree plot) have been proposed (Jolliffe 2002) and a recent mathematically refined approach is the method of the optimal hard threshold for singular values introduced by Gavish and Donoho (2014).

Note that the analysis of variables, which are measured in different units, can be misleading and cause undesirable interpretations. This is because the eigenvectors are not scale invariant. Thus, it is favorable to use the correlation matrix instead of the covariance matrix in general.

[4] Technically, the eigenvectors can be seen as direction cosines, while the corresponding eigenvalues describe the magnitude. By construction, the loading matrix aggregates both.

The singular value decomposition to compute PCA. In practice the explicit computation of the covariance or correlation matrix for massive datasets can be expensive. A more computationally efficient approach to compute the principal components is the singular value decomposition, which avoids the (often) costly computation of the Gram matrix $\mathbf{X}^\top\mathbf{X}$. As described in Section 2.1, the eigenvalue decomposition of the inner and outer dot product of $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$ can be related to the singular value decomposition as follows

$$\mathbf{X}^\top\mathbf{X} = (\mathbf{V}\mathbf{S}\mathbf{U}^\top)(\mathbf{U}\mathbf{S}\mathbf{V}^\top) = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top \tag{15a}$$

$$\mathbf{X}\mathbf{X}^\top = (\mathbf{U}\mathbf{S}\mathbf{V}^\top)(\mathbf{V}\mathbf{S}\mathbf{U}^\top) = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top \tag{15b}$$

where the eigenvalues $\mathbf{\Sigma} = \mathbf{S^2}$ are the squared singular values. The eigenvectors of $\mathbf{X}\mathbf{X}^\top$ are given by the left singular vectors $\mathbf{U}$ and the eigenvectors of $\mathbf{X}^\top\mathbf{X}$ are given by the right singular vectors $\mathbf{V}$ of the matrix $\mathbf{X}$. Thus, from the SVD of $\mathbf{X}$ we recover the rotation matrix $\mathbf{W}$ and eigenvalues $\mathbf{\Lambda}$ of the covariance matrix $\mathbf{C}$ as $\mathbf{\Lambda} = \frac{\mathbf{1}}{\mathbf{n-1}}\mathbf{S^2}$ and $\mathbf{W} = \mathbf{V}$. Moreover, the principal component scores can be computed as

$$\mathbf{Z} = \mathbf{X}\mathbf{W} = \mathbf{U}\mathbf{S}\mathbf{V}^\top\mathbf{W} = \mathbf{U}\mathbf{S}. \tag{16}$$

2.3. Robust principal component analysis.

In the previous two sections we have summarized two closely related methods for dimensionality reduction, namely the singular value decomposition and principal component analysis. In fact, it is well known, as described in Section 2.2, that PCA can be efficiently computed using the SVD. In Section 4, we show that these methods can be used for noise reduction/removal, if the data are corrupted by independent identically distributed (i.i.d.) Gaussian noise and the signal-to-noise ratio is sufficiently large. However, in many practical applications we face data with arbitrarily corrupted observations, e.g., shadows in images or moving objects in videos. In this case the resulting low-rank approximations can be poor, even if only a very few observations are corrupted. Figure 5 presents a toy example of such a possible situation. Subplot (a) shows a grossly corrupted data matrix as the superposition of a half torus (b) and spike noise (c). Specifically, the corrupted data matrix here is superimposed of a low-rank component ($r = 3$) and a sparse component. The question arises, whether a corrupted data matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ can be robustly separated into a low-rank matrix $\mathbf{L} \in \mathbb{R}^{n\times m}$ and sparse matrix $\mathbf{S} \in \mathbb{R}^{n\times m}$

$$\mathbf{A} = \mathbf{L} + \mathbf{S} \tag{17}$$

so that $\mathbf{S}$ captures the perturbations. Indeed, Candès, Li, Ma, and Wright (2011) proved that it is possible to exactly separate such a data matrix into both its low-rank and sparse components, under rather broad assumptions. This is achieved by solving a convenient convex optimization problem, called principal component pursuit (PCP). The objective is to minimize a weighted combination of the nuclear norm $\|\cdot\|_* := \sum_i \sigma_i$ (sum of the singular values) and and the ℓ_1 norm $\|\cdot\|_1 := \sum_{ij} |m_{ij}|$ as follows

$$\min_{\mathbf{L},\mathbf{S}} \|\mathbf{L}\|_* + \lambda\|\mathbf{S}\|_1 \quad \text{subject to } \mathbf{A} - \mathbf{L} - \mathbf{S} = 0. \tag{18}$$

The arbitrary balance parameter, $\lambda > 0$, is typically chosen to be $\lambda = 1/\sqrt{\max(n, m)}$. Hence, the key to achieving such a matrix decomposition is ℓ_1 optimization. The PCP concept is

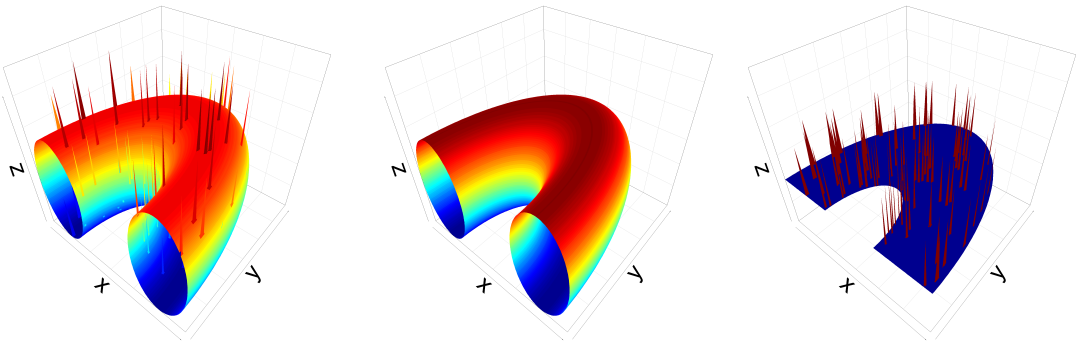

(a) Grossly corrupted torus. (b) Original half torus. (c) Sparse (spiked) noise.

Figure 5: Toy example of a grossly corrupted data matrix. Subplot (a) shows the perturbed torus as a superposition of the half torus (b) and spike noise (c).

mathematically sound and has been applied successfully to images and videos (Wright, Ganesh, Rao, Peng, and Ma 2009). More generally the concept of robustly separating corrupted matrices is denoted as robust principal component analysis (RPCA). However, its biggest challenge is computational efficiency, especially given the iterative nature of the optimization required. The inexact augmented Lagrange multiplier method (Lin *et al.* 2010) is a popular choice, due to its favorable computational properties. This method essentially formulates the following Lagrangian function

$$\mathcal{L}(\mathbf{L}, \mathbf{S}, \mathbf{Y}, \mu) = \|\mathbf{L}\|_* + \lambda\|\mathbf{S}\|_1 + \langle \mathbf{Y}, \mathbf{A} - \mathbf{L} - \mathbf{S}\rangle + \frac{\mu}{2}\|\mathbf{A} - \mathbf{L} - \mathbf{S}\|_F^2 \tag{19}$$

where μ is a positive scalar and the matrix $\mathbf{Y}$ is the Lagrange multiplier. The method achieves an exact separation of the corrupted half torus into the low-rank and sparse components, shown in Figure 6. The reconstruction error is nearly negligible, while in comparison the low-rank SVD approximation contains substantial defects. Its remarkable ability to separate high-dimensional matrices into low-rank and sparse component makes robust principal component analysis an invaluable tool for data science.

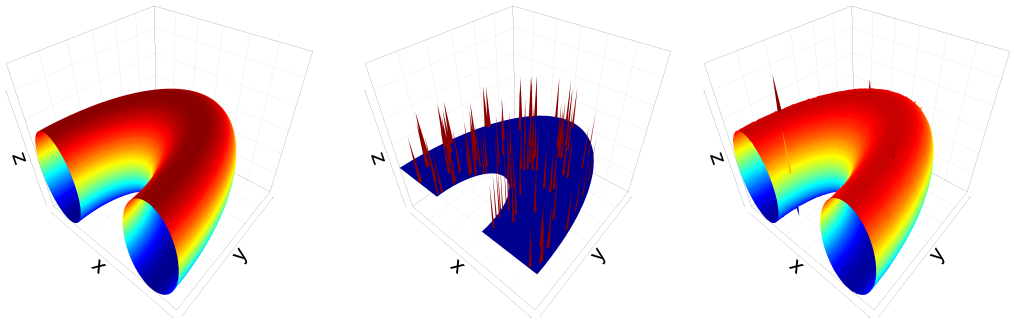

(a) Low-rank component $\mathbf{L}$. (b) Sparse component $\mathbf{S}$. (c) SVD approximation.

Figure 6: Subplot (a) and (b) show the separation of the grossly corrupted torus in Figure 5 using robust PCA. The low-rank component captures the original torus faithfully, and the reconstruction error is as low as 0.0003%. In contrast, the approximation performance of the the low-rank SVD is poor, shown in subplot (c). The reconstruction error is about 1.81%.

3. Randomized algorithms

The computational cost of computing the truncated SVD $\mathbf{A}_k$ using a deterministic algorithm can be tremendous for massive datasets. That is because the cost of a full SVD of an $m \times n$ matrix is of the order $O(mn^2)$, from which the first k components can then be extracted to form $\mathbf{A}_k$. In contrast, randomized matrix techniques can be used to obtain an approximate rank-k singular value decomposition at a cost of $O(mnk)$, substantially smaller than for a full SVD when the dimensions of $\mathbf{A}$ are large. The randomized low-rank SVD approach can also be applied to accelerate the PCA computation. This section shows how this is accomplished.

3.1. Randomized singular value decomposition.

We present details of the randomized low-rank SVD algorithm which comes with favorable error bounds relative to the optimal truncated SVD, as presented in the seminal paper by Halko *et al.* (2011b) and further analyzed and implemented in Voronin and Martinsson (2015).

The algorithm can be conceptually divided into two stages. Given a real matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ and a desired target rank $k \ll m, n$, the first stage makes use of randomization to build a low-dimensional subspace that approximately captures the column space of $\mathbf{A}$. This is achieved by simply drawing k random vectors $\omega_1, \ldots, \omega_k$ from a sub-Gaussian distribution, followed by computing a random projection of the data matrix

$$\mathbf{y}_i = \mathbf{A}\omega_i \quad \text{for } i\text{=1,2,....,}k\ . \tag{20}$$

The set of random projections $\{\mathbf{y}_i\}$ serve as the approximate basis for the column space of the data matrix. This is because probability theory guarantees that random vectors, and hence the corresponding random projections, are linearly independent with high probability and quickly sample the range of a numerical rank deficient matrix $\mathbf{A}$. More compactly, the operation in (20) can be performed via one matrix–matrix multiplication:

$$\mathbf{Y} = \mathbf{A}\boldsymbol{\Omega} \tag{21}$$

where $\mathbf{Y} \in \mathbb{R}^{m\times k}$ is the sample matrix and $\boldsymbol{\Omega} \in \mathbb{R}^{n\times k}$ is a random sub-Gaussian matrix.

A natural basis can then be obtained by computing the economic QR decomposition of $\mathbf{Y} = \mathbf{QR}$, where $\mathbf{Q} \in \mathbb{R}^{m\times k}$ is orthonormal and has the same column space as $\mathbf{Y}$. If k is large enough and close to the numeric rank $\bar{r}$ of $\mathbf{A}$, then the resulting approximate basis $\mathbf{Q}$ satisfies

$$\mathbf{A} \approx \mathbf{QQ}^\top \mathbf{A}. \tag{22}$$

Notice that the statement in (22) is replaced by an equality when the number of samples drawn is large enough so that $R(\mathbf{A}) \subseteq R(\mathbf{Y})$. The second stage of the algorithm is concerned with obtaining the low-rank singular value decomposition from a small compressed matrix, which faithfully captures the relevant spectral information of the original data matrix. Therefore, we first project the input matrix $\mathbf{A}$ onto the low-dimensional subspace

$$\mathbf{B} = \mathbf{Q}^\top \mathbf{A} \tag{23}$$

to obtain the relatively small (if $k \ll m, n$) matrix $\mathbf{B} \in \mathbb{R}^{k\times n}$. Computing any deterministic matrix factorization of this small compressed matrix is relatively inexpensive (if n is very large,

it is possible to work with $\mathbf{B}\mathbf{B}^\top$ instead of directly with $\mathbf{B}$). Hence one obtains the following approximate SVD by using a standard deterministic algorithm

$$\mathbf{B} = \tilde{\mathbf{U}}\mathbf{\Sigma}\mathbf{V}^\top. \tag{24}$$

It then remains to recover the left singular vectors as follows

$$\mathbf{U} \approx \mathbf{Q}\tilde{\mathbf{U}}. \tag{25}$$

The justification for the randomized algorithm can be sketched as follows starting with (22):

$$\begin{aligned} \mathbf{A} &\approx \mathbf{Q}\mathbf{Q}^\top\mathbf{A} \\ &\approx \mathbf{Q}\mathbf{B} \\ &\approx \mathbf{Q}\tilde{\mathbf{U}}\mathbf{\Sigma}\mathbf{V}^\top \\ &\approx \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top. \end{aligned} \tag{26}$$

Figure 7 illustrates the computational steps involved to compute the randomized SVD.

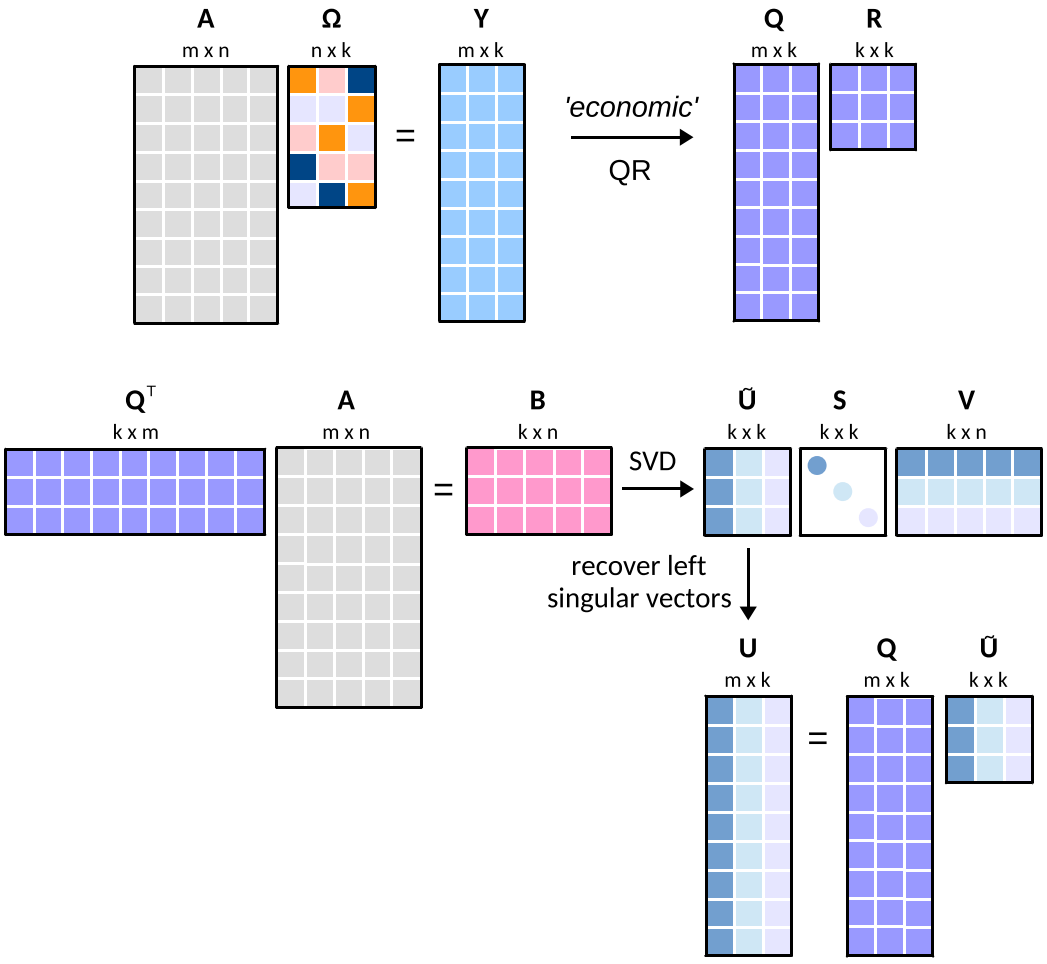

Figure 7: Schematic architecture of the randomized singular value decomposition.

3.2. Computational considerations.

The approximation error of the randomized SVD can be controlled by two tuning parameters. First, it is beneficial to introduce a small oversampling parameter p, i.e., instead of using k random vectors, $k+p$ random vectors are used to obtain the sample matrix $\mathbf{Y}$. Often, a fairly small oversampling parameter is sufficient, e.g., $p = 5, 10$. In Halko *et al.* (2011b), the following practical error bound is derived:

$$\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^\top \mathbf{A}\| \leq \left[1 + 11\sqrt{k+p}\sqrt{\min(m,n)}\right] \sigma_{k+1} \tag{27}$$

where we compare the right hand side to σ_{k+1} in (4). The above bound holds with probability $1 - 3p^{-p}$ for $p \geq 4$, which justifies the use of a small oversampling parameter p relative to k.

The second method to improve the approximation involves the use of power sampling iterations. Instead of obtaining the sampling matrix $\mathbf{Y}$ directly from the data matrix, one samples from a pre-processed matrix as follows

$$\mathbf{Y} = \big((\mathbf{A}\mathbf{A}^\top)^q \mathbf{A}\big)\boldsymbol{\Omega} \tag{28}$$

where q is an integer specifying the number of power iterations. A simple computation shows that $(\mathbf{A}\mathbf{A}^\top)^q \mathbf{A} = \mathbf{U}\mathbf{S}^{2q+1}\mathbf{V}$. Hence, for $q > 0$, the modified sampling matrix has a relatively fast decay of singular values even when the decay in $\mathbf{A}$ is modest. The drawback is that additional matrix-matrix multiplications are required. However, when the singular values of the data matrix decay slowly, about $q = 1, 2$ power iterations can considerably improve the approximation. In a practical implementation subspace iterations are used instead of power iterations for numerical stability (Gu 2015). The algorithm in Figure 8 summarizes an implementation with an oversampling parameter and subspace iterations.

function $[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \texttt{rsvd}(\mathbf{A}, k, p, q)$

(1) $l = k + p$ — slight oversampling
(2) $\boldsymbol{\Omega} = \texttt{rnorm}(n, l)$ — generate sampling matrix
(3) $\mathbf{Y} = \mathbf{A}\boldsymbol{\Omega}$ — draw range samples
(4) **for** $j = 1, \ldots, q$ — perform optional subspace iterations
(5) $\quad[\mathbf{Q}, \sim] = \texttt{qr}(\mathbf{Y})$
(6) $\quad\mathbf{Z} = \mathbf{A}^\top \mathbf{Q}$
(7) $\quad[\mathbf{Q}, \sim] = \texttt{qr}(\mathbf{Z})$
(8) $\quad\mathbf{Y} = \mathbf{A}\mathbf{Q}$
(9) **end for**
(10) $[\mathbf{Q}, \sim] = \texttt{qr}(\mathbf{Y})$ — form orthonormal samples matrix
(11) $\mathbf{B} = \mathbf{Q}^\top \mathbf{A}$ — project to smaller space
(12) $[\tilde{\mathbf{U}}, \mathbf{S}, \mathbf{V}] = \texttt{svd}(\mathbf{B})$ — compact SVD of smaller matrix B
(13) $\mathbf{U} = \mathbf{Q}\tilde{\mathbf{U}}$ — recover left singular vectors
(14) $\mathbf{U} = \mathbf{U}(:, 1:k),\ \mathbf{s} = \mathbf{S}(1:k, 1:k),\ \mathbf{V} = \mathbf{V}(:, 1:k)$ — extract k components

Figure 8: A randomized SVD algorithm.

Remark 1. The Gaussian or uniform distribution are common choices to generate $\mathbf{\Omega}$.

Remark 2. Numerical examples in Section 4 show that good default values for both the oversampling parameter and the number of subspace iterations are $p = 10$ and $q = 1$.

Remark 3. The computational time can be further reduced by first computing the QR decomposition of $\mathbf{B}^\top$ and then computing the SVD of the even smaller matrix $\mathbf{R} \in \mathbb{R}^{l \times l}$ or via the eigendecomposition of the $l \times l$ matrix $\mathbf{B}\mathbf{B}^\top$ (see Voronin and Martinsson (2015) and Section 3.4 for further details).

Error bounds. The randomized algorithm for the low-rank SVD only approximates $\mathbf{A}_k$, the actual rank-k truncated SVD, obtained from the full SVD of $\mathbf{A}$. The error bounds are worse than the optimal ones given in (4), but the practical results tend to be very close to optimal. In the algorithm in Figure 8, we compute the full SVD of $\mathbf{B}$, so it follows that $\|\mathbf{A} - \mathbf{U}_k\mathbf{S}_k\mathbf{V}_k^\top\| = \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^\top\mathbf{A}\|$. In Halko *et al.* (2011b), the following error bound is derived:

$$\mathsf{E}\left[\|\mathbf{A} - \mathbf{U}_k\mathbf{S}_k\mathbf{V}_k^\top\|\right] \leq \left[\left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1}^{2q+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{j>k}\sigma_j^{2(2q+1)}\right)^{\frac{1}{2}}\right]^{\frac{1}{2q+1}}$$

When the singular values σ_j for $j > k$ are small in magnitude, the sum in the above formula is small and the bound approaches the theoretically optimal value of σ_{k+1} with increasing q.

Note again that the error in the appoximate low rank SVD factorization obtained with the randomized algorithm coincides with the bound $\|\mathbf{A} - \mathbf{Q}\mathbf{Q}^\top\mathbf{A}\|$. To control the error, one may proceed to buildup $\mathbf{Q}$ by adding samples to expand $\mathbf{Y}$ a step at a time, until the bound becomes small. A more intricate iterative procedure for building up $\mathbf{Q}$ such that the bound holds to a specified tolerance is presented in Martinsson and Voronin (2015).

3.3. Measurement matrices.

An essential computational step of the randomized algorithm is to construct a measurement matrix $\mathbf{\Omega}$ that is used to sample the range of the input matrix. Therefore we rely on the properties of random vectors, which imply that the randomly generated columns of the measurement matrix are linearly independent with high probability. Indeed, this includes the whole class of sub-Gaussian random variables (Rivasplata 2012), including matrices with Bernoulli or uniform random entries.

Due to its attractive theoretical properties, the random Gaussian measurement matrix is the most prominent choice, consisting of independent identically distributed (i.i.d.) $\mathcal{N}(0, 1)$ standard normal entries. In practice, however, uniform random measurements are sufficient and are less expensive to generate. The drawback is that dense vector and matrix operations are expensive for large matrices. However, BLAS (Basic Linear Algebra Subprograms) operations tend to be highly scalable and computations could be substantially accelerated with parallel or distributed processing, e.g., a graphics processing unit (GPU) implementation (Erichson and Donovan 2016; Voronin and Martinsson 2015).

Woolfe *et al.* (2008) proposed another computationally efficient approach, exploiting the properties of a structured random measurement matrix. For instance, using a subsampled random Fourier transform measurement matrix reduces the computational costs from $O(mnk)$ to $O(mn\log(k))$.

3.4. Randomized principal component analysis.

The information in the data is often explained by just the first few dominant principal components, and thus the randomized singular value decomposition provides an efficient and fast algorithm for computing the first k principal components. This approach is denoted randomized principal component analysis (rPCA) as introduced by Rokhlin, Szlam, and Tygert (2009) and Halko, Martinsson, Shkolnisky, and Tygert (2011a). See also Szlam, Kluger, and Tygert (2014) for some interesting implementation details. Instead of using the randomized SVD Algorithm in Figure 8 directly, we use a modified Algorithm proposed by Voronin and Martinsson (2015). Hence, instead of computing the SVD of $\mathbf{B}$, the eigendecomposition of the smaller $\mathbf{B}\mathbf{B}^\top$ matrix is computed.

Specifically, for use in PCA, we compute the approximate randomized low-rank eigenvalue decomposition of $\mathbf{X}^\top\mathbf{X}$ given the rectangular mean centered matrix $\mathbf{X}$. Note that if $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, then $\mathbf{X}^\top\mathbf{X} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^\top \approx \mathbf{V}_k\mathbf{\Sigma}_k^2\mathbf{V}_k^\top$ and $\mathbf{X}\mathbf{A}^\top = \mathbf{U}\mathbf{\Sigma}^2\mathbf{U}^\top \approx \mathbf{U}_k\mathbf{\Sigma}_k^2\mathbf{U}_k^\top$. Performing such an eigendecomposition for an $m \times n$ matrix with many columns would be expensive, since $\mathbf{X}^\top\mathbf{X}$ is $n \times n$. On the other hand, the first k singular vectors and values of $\mathbf{A}$ are approximately captured by $\mathbf{B} = \mathbf{Q}^\top\mathbf{X}$ where $\mathbf{Q}$ is such that $\mathbf{Q}\mathbf{Q}^\top\mathbf{X} \approx \mathbf{X}$. The $l \times n$ matrix $\mathbf{B}$ could still be large, so instead we can work with the $l \times l$ matrix $\mathbf{B}\mathbf{B}^\top$. Notice that if $\mathbf{B} = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}\tilde{\mathbf{V}}$, which we do not compute, then $\mathbf{B}\mathbf{B}^\top = \tilde{\mathbf{U}}\tilde{\mathbf{\Sigma}}^2\tilde{\mathbf{U}}^\top$. Suppose that we instead construct the eigendecomposition of $\mathbf{B}\mathbf{B}^\top = \tilde{\mathbf{W}}\tilde{\mathbf{S}}\tilde{\mathbf{W}}^\top$. It follows that $\tilde{\mathbf{\Sigma}} = \tilde{\mathbf{S}}^{\frac{1}{2}}$ and that $\text{trunc}(\tilde{\mathbf{\Sigma}}, k) \approx \text{trunc}(\mathbf{\Sigma}, k)$. The first k left singular vectors can be recovered via the computation $\mathbf{U}_k = \text{trunc}(\mathbf{Q}\tilde{\mathbf{W}}, k)$. Employing the expression from Section 2.1, we find that the first k right eigenvectors are approximately recovered via $\mathbf{V}_k = \text{trunc}(\mathbf{B}^\top\tilde{\mathbf{W}}\tilde{\mathbf{\Sigma}}^{-1}, k) = \text{trunc}(\mathbf{B}^\top\tilde{\mathbf{W}}\tilde{\mathbf{S}}^{-\frac{1}{2}}, k)$. Note that in practice, if only the eigendecomposition of $\mathbf{X}^\top\mathbf{X}$ or $\mathbf{X}\mathbf{X}^\top$ is needed, then only one of $\mathbf{V}_k$ or $\mathbf{U}_k$ needs to be computed.

The procedure is summarized in Algorithm 9 and the key differences are in lines $12 - 15$.

function $[\mathbf{\Lambda}, \mathbf{V}] = \texttt{reigen}(\mathbf{X}, k, p, q)$

(1)	$l = k + p$	slight oversampling
(2)	$\mathbf{\Omega} = \texttt{rnorm}(n, l)$	generate sampling matrix
(3)	$\mathbf{Y} = \mathbf{X}\mathbf{\Omega}$	draw range samples
(4)	**for** $j = 1, \ldots, q$	perform optional subspace iterations
(5)	$[\mathbf{Q}, \sim] = \texttt{qr}(\mathbf{Y})$	
(6)	$\mathbf{Z} = \mathbf{X}^\top\mathbf{Q}$	
(7)	$[\mathbf{Q}, \sim] = \texttt{qr}(\mathbf{Z})$	
(8)	$\mathbf{Y} = \mathbf{X}\mathbf{Q}$	
(9)	**end for**	
(10)	$[\mathbf{Q}, \sim] = \texttt{qr}(\mathbf{Y})$	form orthonormal samples matrix
(11)	$\mathbf{B} = \mathbf{Q}^\top\mathbf{X}$	project to smaller space
(12)	$\mathbf{BBt} \leftarrow \mathbf{B} * \mathbf{B}^\top$	compute outer product
(13)	$[\mathbf{\Lambda}, \tilde{\mathbf{W}}] = \texttt{eigen}(\mathbf{BBt})$	compact eigendecomposition
(14)	$\mathbf{V} = \mathbf{B}^\top * \tilde{\mathbf{W}} * \mathbf{\Lambda}^{-0.5}$	recover right eigenvectors
(15)	$\mathbf{V} = \mathbf{V}(:, 1:k),\ \mathbf{\Lambda} = \mathbf{\Lambda}(1:k, 1:k)$	extract k components

Figure 9: A randomized eigendecomposition algorithm.

4. Numerical Examples

In this section, we demonstrate the usage of randomized matrix algorithms in practice. We illustrate the performance of our routines using several standard examples and compare the results to the corresponding deterministic R functions. Section 4.1 starts with a classic example showing how the randomized singular value decomposition function `rsvd()` can be used for image compression. Sections 4.2 and 4.3 present two examples illustrating the randomized principal component analysis routine `rpca()`. The latter example involves a large dense matrix and highlights the computational advantage of randomized PCA. Section 4.4 introduces the randomized robust principal component analysis function `rrpca()`. Finally, Section 4.5 investigates the performance of the `rsvd()` function, showing speed-ups ranging from 5 to 150 times. In the following it is assumed that the **rsvd** package is installed and loaded.

```
R> install.packages("rsvd")
R> library("rsvd")
```

A workstation with an Intel Xeon CPU E5-2620 2.4GHz, 32GB DDR3 memory, and the operating-system Ubuntu 16.04 LTS is used for all following computations and the **microbenchmark** package is utilized for evaluating the computational runtime of the algorithms (Mersmann, Beleites, Hurling, and Friedman 2015).

4.1. SVD example: Image compression.

The singular value decomposition can be used to obtain a low-rank approximation to high-dimensional data. Image compression is a simple and illustrative example of this.[5] Although images often feature a high-dimensional ambient space, the underlying structure can be represented by a very sparse model. This means that most natural images can be faithfully recovered from a relatively small set of basis functions. For demonstration a 1600×1200 grayscale image is used in the following.

```
R> data("tiger")
R> image(tiger, col = gray(0:255 / 255))
```

A grayscale image may be thought of as a real-valued matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, where m and n are the number of pixels in the vertical and horizontal directions, respectively. To compress the image we must first decompose the image. The singular vectors and values provide a hierarchical representation of the image in terms of a new coordinate system defined by dominant correlations within the image. Thus, the number of singular vectors used for approximation poses a trade-off between the compression rate (i.e., the number of singular vectors to be stored) and the image details. First, the R base `svd()` function is used to compute the singular value decomposition

```
R> k <- 100
R> tiger.svd <- svd(tiger, nu = k, nv = k)
```

The `svd()` function returns three objects: `u`, `v` and `d`. The first two objects are $m \times k$ and $n \times k$ arrays, namely the truncated right and left singular vectors. The vector `d` comprises the

[5] Among the many strategies to compress or denoise images, the singular value decomposition is one prominent tool, although it is certainly not the most effective.

$min(m, n)$ singular values in descending order. Now, the dominant $k = 100$ singular values are retained to approximate/reconstruct ($\mathbf{A}_k = \mathbf{U}_k\mathbf{D}_k\mathbf{V}_k^\top$) the original image

```
R> tiger.re <- tiger.svd$u %*% diag(tiger.svd$d[1:k]) %*% t(tiger.svd$v)
R> image(tiger.re, col = gray(0:255 / 255))
```

The normalized root mean squared error (nrmse) is a common measure for the reconstruction quality of images (Fienup 1997), computed as

```
R> nrmse <- sqrt(sum((tiger - tiger.re) ** 2 ) / sum(tiger ** 2))
```

Using only the first $k = 100$ singular values/vectors, a reconstruction error as low as 0.121 is achieved. This illustrates that in general natural images feature a very compact representation.

The `svd()` function provides an interface to the underlying LAPACK SVD routines (Anderson *et al.* 1999). These routines are known to be numerical stable and highly accurate, but computationally demanding. Specifically, the computational time required to approximate large-scale data is tremendous. However, if the data matrix exhibits a low-rank structure, the randomized algorithm can ease the computational time substantially. The provided `rsvd()` function can be used as a plug-in function for the base `svd()` function, in order to compute the near-optimal low-rank singular value decomposition

```
R> tiger.rsvd <- rsvd(tiger, k = k)
```

Similar to the base SVD function, the `rsvd()` function returns three objects: `u`, `v` and `d`. Again, `u` and `v` are $m \times k$ and $n \times k$ arrays containing the approximate right and left singular vectors and the vector `d` comprises the k singular values in descending order. The approximation accuracy of the rSVD algorithm can be controlled by two parameters p and q. The former is an oversampling parameter and the latter controls the number of subspace iterations. Specifically, if the singular value spectrum decays slowly, an increased number of subspace iterations can improve the accuracy. However, in practice we have not encountered a situation which requires more then three subspace iterations ($q > 3$). By default the parameters are set to $p = 10$ and $q = 1$. Again, the approximated image and its reconstruction quality can be obtained as

```
R> tiger.re <- tiger.rsvd$u %*% diag(tiger.rsvd$d) %*% t(tiger.rsvd$v)
R> nrmse <- sqrt(sum((tiger - tiger.re) ** 2) / sum(tiger ** 2))
```

The reconstruction quality is almost similar, with an insignificantly larger nrmse of about 0.125, however, the randomized algorithm is substantially faster. Figure 10 presents the results of the image approximations using both the deterministic and randomized SVD algorithm. By visual inspection no significant difference can be seen between (b) and (d). However, it can be seen that the quality slightly suffers by excluding the computation of at least one subspace iteration (c).

Figure 11 shows the corresponding singular values and it can be seen that the randomized algorithm faithfully captures the true singular values with $q = \{1, 2\}$ subspace iterations. However, without computing subspace iterations, the singular values decay slowly. Table 1 shows the median elapsed run-time for different svd algorithms in R. The `rsvd()` function achieves an average speed-up of about 7 over the `svd()` function. The `svds()` function of

(a) Original image.

(b) SVD. (nrmse=0.121)

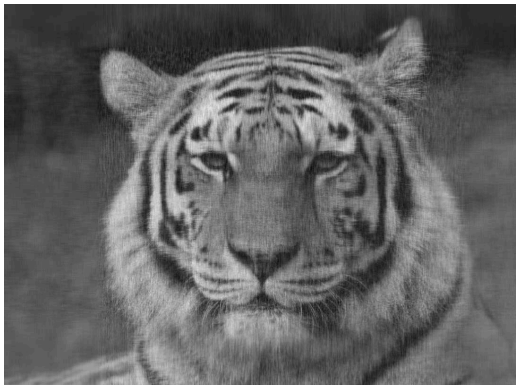

(c) rSVD using $q = 0$. (nrmse=0.163)

(d) rSVD using $q = 1$. (nrmse=0.125)

Figure 10: Image compression using the SVD. Subplot a) shows the original image and subplots (b), (c) and (d) show the reconstructed image using the dominant $k = 100$ singular vectors by both the deterministic and randomized SVD algorithm. The reconstruction quality using $q = 1$ subspace iterations is insignificant compared to the deterministic SVD algorithm.

Method	**Parameters**	**Time (s)**	**Speedup**	**Error**
`svd()`		1.05	-	0.121
`propack.svd()`	`neig=100`	0.94	1.11	0.121
`svds()`	`k=100`	0.25	4.21	0.121
`rsvd()`	`k=100, q=0`	0.08	12.80	0.163
	`k=100, q=1`	0.15	7.05	0.125
	`k=100, q=2`	0.22	4.86	0.122

Table 1: Summary of algorithms runtime and errors. The randomized routines achieve a substantial speedup, while attaining similar reconstruction errors with $q \geq 1$.

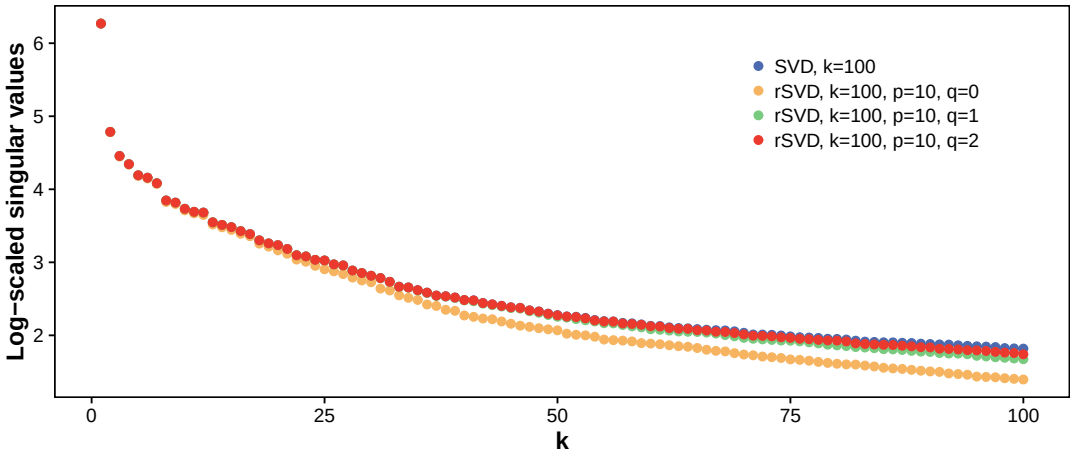

Figure 11: Dominant log-scaled singular value spectrum of the image 'tiger'.

the **Rspectra** package achieves a speedup of about 4.2. However, the computational gain of the randomized algorithm becomes more pronounced with increased matrix dimension, e.g., images with higher resolution.

4.2. PCA example: Edgar Anderson's Iris Data.

The `princomp()` and `prcomp()` functions are the standard methods to compute the principal component analysis in R. The former method is based on the eigenvalue decomposition on the correlation or covariance matrix, while the latter method is computing the singular value decomposition of the centered and possibly scaled data matrix (Venables and Ripley 2002). The SVD based method is the preferred approach in practice due to its better numerical stability. The function `rpca()` is designed as a substitute for `prcomp()` function. The key difference is that a randomized algorithm is used to perform the underlying computations. By default, the choice of whether to use the randomized or base SVD algorithm is made automatically. If the target-rank $k > n/1.5$ the randomized algorithm is in general inefficient, and the deterministic SVD algorithm is favored.

In the following we use the famous `iris` data set, collected by Anderson (1935), to demonstrate the usage of the `rpca()` function. The data set comprises 50 observations from each of 3 species of iris: 'setosa', 'versicolor', and 'virginica'. Each observation is a vector of values, measured in centimeters, of the variables 'sepal length', 'sepal width', 'petal length', and 'petal width', respectively.

```
R> data("iris")
R> head(iris, 2)
```

```
  Sepal.Length  Sepal.Width  Petal.Length  Petal.Width  Species
1          5.1          3.5           1.4          0.2   setosa
2          4.9          3.0           1.4          0.2   setosa
```

Following Venables and Ripley (2002), first the log transformed measurements are computed

```
R> log.iris <- log( iris[ , 1:4] )
```

which is a standard skewness transformation. Our primary aim here is to visualize the data in 2-D space using the first two ($k = 2$) principal components.

```
R> iris.rpca <- rpca(log.iris, k = 2)
```

By default, the data are mean centered and standardized, i.e., `center = TRUE` and `scale = TRUE`.[6] This means that the correlation matrix is implicitly computed here. Set `scale = FALSE` to use the implicit covariance matrix instead. The returned object of the `rpca()` function has the same summary, print and plot capabilities as `prcomp()`. For example the summary function shows

```
R> summary(iris.rpca)
```

```
                         PC1   PC2
Explained variance     2.933 0.907
Standard deviations    1.712 0.952
Proportion of variance 0.733 0.227
Cumulative proportion  0.733 0.960
Eigenvalues            2.933 0.907
```

Thus, about 73% of the variability in the data is explained by just the first PC and about 23% by the second. Note, that the explained variance of the principal components corresponds exactly to the eigenvalues. In addition, the print function shows the variable loadings on the PCs. For instance, the first PC is best explained by 'sepal length', 'petal length' and 'petal width', while 'sepal width' loads onto the second component.

```
R> print(iris.rpca)
```

```
Standard deviations:
[1] 1.712 0.952

Eigenvalues:
[1] 2.933 0.907

Rotation:
               PC1    PC2
Sepal.Length  0.504 -0.455
Sepal.Width  -0.302 -0.889
Petal.Length  0.577 -0.034
Petal.Width   0.567 -0.035
```

Two widely used methods for visualizing the results of the principal component analysis are the correlation and bi-plot. The correlation plot for the first and second principal component can be produced as follows

[6] Note, that the `prcomp()` function has by default the arguments `center = TRUE` and `scale = FALSE`.

```
R> ggcorplot(iris.rpca, pcs = c(1, 2))
```

In order to create the biplot, the principal component scores (rotated variables) must be computed. This can be achieved by passing the argument `retx=TRUE` to the `rpca()` function. However, this requires re-computation and instead we use the provided predict function

```
R> iris.rpca$x <- predict(iris.rpca, log.iris)
```

Then the biplot is obtained as

```
R> ggbiplot(iris.rpca, groups = iris$Species)
```

Figure 12 shows the two plots, based on the **ggplot2** package (Wickham 2009). The correlation plot on the left visualizes the original variables by their correlation with the principal components (Abdi and Williams 2010). It can be seen that the variable 'sepal width' is most strongly correlated with the first two PCs. Moreover, it reveals the inter-relationships between the original variables. Clearly, the variables 'petal length' and 'petal width' are highly correlated, i.e., they carry redundant information. Whereas, 'sepal width' is only slightly correlated with the other variables. The correlations can be also computed as row sums of the squared eigenvectors, e.g., `rowSums(iris.rpca$rotation**2)`. The biplot represents both the observations and variables in eigenspace, i.e., it is a scatter plot of the principal component scores overlaid with the variable loadings. In particular, the biplot may help to reveal underlying patterns. For instance, here 'setosa' is distinct from the two other species. Further, the first PC is mainly explained by 'sepal length', 'petal length' and 'petal width', as noted before. An excellent general reference for biplots is Greenacre (2010).

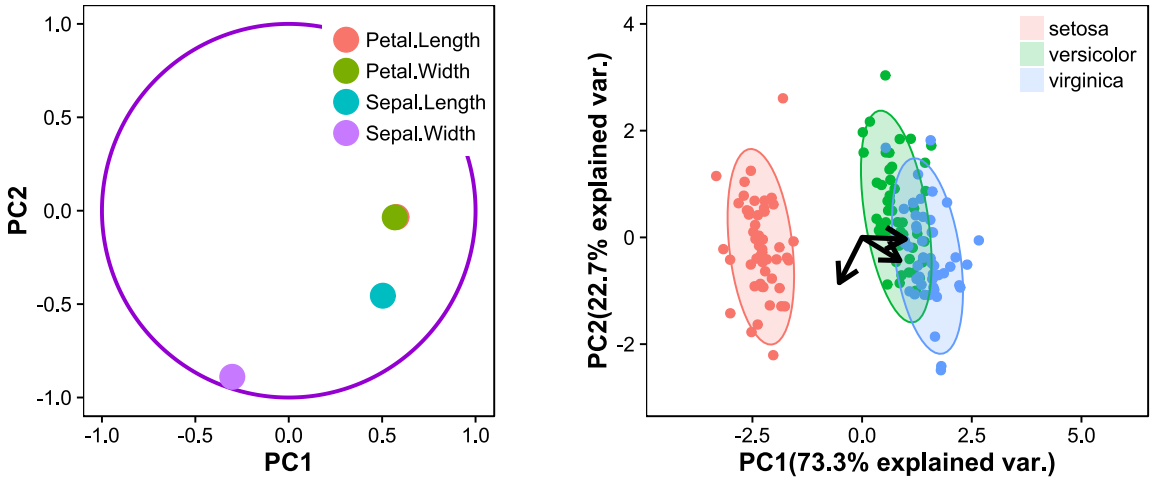

Figure 12: The left subplot visualizes the correlations of the original variable with the PCs. The biplot on the right plots the principal component scores, overlaid with a plot showing the principal directions.

4.3. PCA example: Eigenfaces.

One of the most striking demonstrations of PCA are eigenfaces, first studied by Kirby and Sirovich (1990). The aim is to extract the most dominant correlations between different faces from a large set of facial images. Specifically, the resulting columns of the rotation matrix (i.e., the eigenvectors) represent 'shadows' of the faces, the so-called eigenfaces. Specifically, the eigenfaces reveal both inner face features (e.g., eyes, nose, mouth) and outer features (e.g., head shape, hairline, eyebrows). These features can then be used for facial recognition and classification as first shown by Turk and Pentland (1991).

In the following we use the downsampled cropped Yale face database B (Georghiades, Belhumeur, and Kriegman 2001). The dataset comprises 2410 grayscale images of 38 different people, cropped and aligned, and can be loaded as

```
R> download.file("https://github.com/Benli11/data/raw/master/R/faces.RData",
+ "faces.RData")
R> load("faces.RData")
```

For computational convenience the 96×84 faces images are stored as column vectors of the data matrix. For instance, the first face can be displayed as

```
R> face <- matrix(rev(faces[ , 1]), nrow = 84, ncol = 96)
R> image(face, col = gray(0:255 / 255))
```

In order to approximate the $k = 20$ dominant eigenfaces, we can proceed as follows

```
R> faces.rpca <- rpca(t(faces), k = 20, center = TRUE, scale = TRUE)
```

Note, that the data matrix needs to be transposed, so that each column corresponds to a pixel location rather then to a person. Here, the analysis is performed on the correlation matrix by setting the argument `scale = TRUE`. The summary is as follows

```
R> summary(faces.rpca)

                            PC1         PC2         PC3     ...
Explained variance     2901.539    2706.699     388.053     ...
Standard deviations      53.866      52.026      19.699     ...
Proportion of variance    0.360       0.336       0.048     ...
Cumulative proportion     0.360       0.695       0.744     ...
Eigenvalues            2901.539    2706.699     388.053     ...
```

Just the first 3 PCs explain about 74% of the total variation in the data, while the first 20 PCs explain more then 88%. The summary can be visualized using either the provided standard `plot()` or the pretty plot functions, as shown in Figure 13.

```
R> ggscreeplot(faces.rpca)
R> ggscreeplot(faces.rpca, type =  "cum")
R> ggscreeplot(faces.rpca, type = "ratio")
```

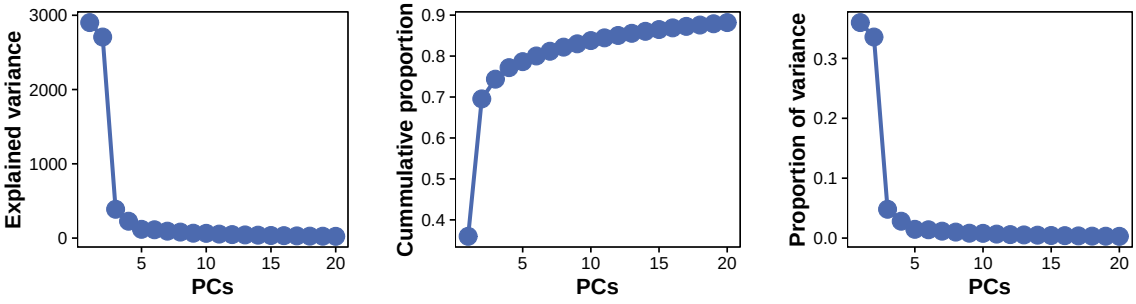

Figure 13: The left plot shows the explained variance (eigenvalues) in decaying order. The middle plot shows the cumulative proportion of the explained variance, and the right plot shows the proportion explained by the principal components.

Finally, the eigenvectors can be visualized as eigenfaces, e.g., the first eigenvector (eigenface) is displayed as follows

```
R> eigenface <- matrix(rev(faces.rpca$rotation[ , 1]), nrow = 84, ncol = 96)
R> image(eigenface, col = gray(0:255 / 255))
```

The mean face is provided as the attribute `center`, and displayed as

```
R> meanface <- matrix(rev(faces.rpca$center), nrow = 84, ncol = 96)
R> image(meanface, col=gray((0:255)/255))
```

Figure 14 shows the first six dominant eigenfaces, computed with both the `prcomp()` and the `rpca()` function. The eigenfaces encode the facial features as well as the illumination. To better contextualize the key features different color maps can be used. In subplot (b) and (d) the colormap `cm.colors(255)` is used, better characterizing the facial shapes (first eigenface), and illumination variations (second eigenface). The third and fifth eigenface feature the nose and eyebrows. Figure 15 shows the standard deviations (eigenvalues) for the `prcomp()` and `rpca()` function. The randomized algorithm faithfully approximates the eigenvalues. Note that the color maps for some of the eigenfaces are flipped. This is because the signs of the columns of the rotation matrix are arbitrary and differ between different PCA implementations. The corresponding computational timings and errors are listed in Table 2. The randomized algorithm achieves a speedup of about 10, while attaining a similar approximation error.

Method	**Parameters**	**Time (s)**	**Speedup**	**Error**
`prcomp()`	-	19.13	-	0.228
`rsvd()`	`k=20, q=1`	1.90	10.05	0.232
	`k=20, q=2`	2.11	9.08	0.229

Table 2: Computational time for the deterministic and randomized SVD algorithm. The randomized routines achieve a substantial speed-up averaged over 200 runs. The randomized algorithms achieves an about 33 fold speed-up compared to the deterministic PCA algorithm.

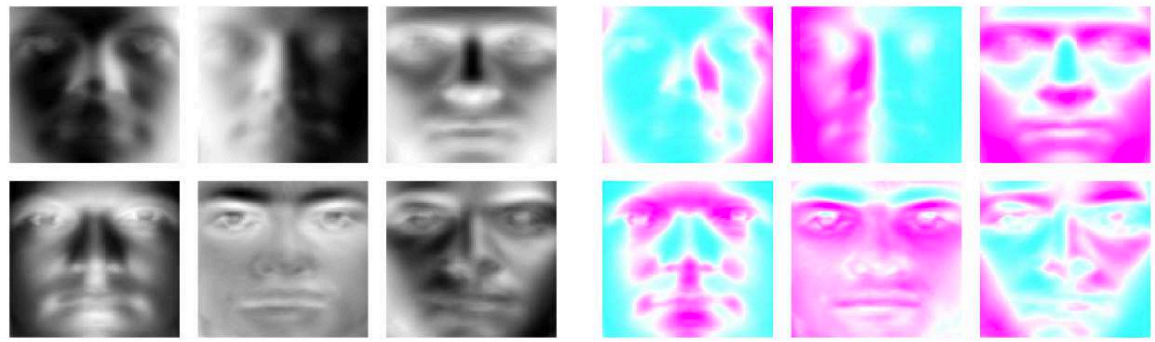

(a) Eigenfaces.

(b) Colored eigenfaces.

(c) Randomized eigenfaces.

(d) Colored randomized eigenfaces.

Figure 14: Dominant six eigenfaces computed with both the `prcomp()` and `rpca()` functions. Subplot (a) and (c) show the graycolor and (b) and (c) the colored eigenfaces using the colormap `cm.colors(255)`. Note that some of the randomized eigenfaces are sign flipped.

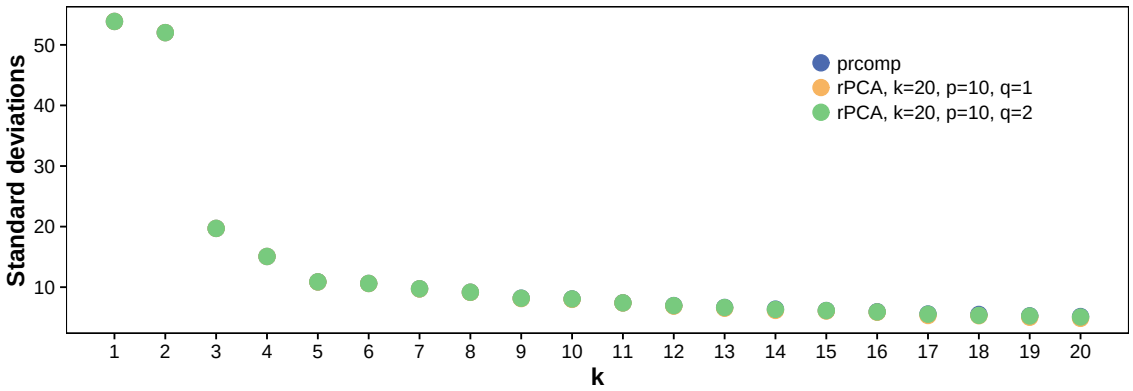

Figure 15: Standard deviations for the deterministic and randomized PCA algorithm.

4.4. Robust PCA example: Foreground/background separation.

In the following we demonstrate the randomized RPCA algorithm for video foreground/background separation, a problem studied widely in the computer vision community (Bouwmans and Zahzah 2014). Background modeling is a fundamental task in computer vision to detect moving objects in a given video stream from a static camera. The key idea is that dynamic pixels in successive video frames are considered foreground objects, whereas static pixels are considered part of the background. Thus the foreground can be found in a video by removing the background. However, background estimation is a challenging task due to the presence of foreground objects and variability in the background itself, e.g., waving trees, water fountains or illumination changes. One way to tackle this challenge is to exploit the low-rank structure of the background, while considering foreground objects as outliers. A solution to this problem is provided by robust principal component analysis, which separates a matrix into a low-rank (background) and sparse component (representing the activity in the scene).[7] An example surveillance video (Goyette, Jodoin, Porikli, Konrad, and Ishwar 2012) containing 200 grayscale frames of size 176×144 can be loaded as follows

```
R> download.file(https://github.com/Benli11/data/raw/master/R/highway.RData",
+ "highway.RData")
R> load("highway.RData")
```

Each frame is stored as a flattened column vector. For instance the 200th frame can be displayed as

```
R> image(matrix(highway[ , 200], ncol = 144 , nrow = 176),
+ col = gray(0:255 / 255))
```

Assume that the background is the low-rank component as discussed above. This assumption is reasonable since the camera is fixed and the background only gradually changes over time. The separation is then obtained as follows

```
R> highway.rrpca <- rrpca(highway, k = 10, p = 5, trace = TRUE)
```

The `rrpca()` function returns two attributes: `L` and `S`, which are both $m \times n$ arrays. The first array contains the low-rank component and the latter the sparse component. Thus, the results can be illustrated by displaying an example frame of both arrays

```
R> image(matrix(highway.rrpca$L[ , 200], ncol = 144 , nrow = 176),
+ col = gray(0:255 / 255))
R> image(matrix(highway.rrpca$S[ , 200], ncol = 144 , nrow = 176),
+ col = gray(0:255 / 255))
```

as shown in Figure 16. It clearly can be seen that the algorithm treats the foreground objects as sparse components, representing outlying entries in the data matrix. Thus the algorithm faithfully separates the video into its two components. Figure 17 shows the convergence

[7] A general drawback of RPCA methods is that they rely on one or more tuning parameters, although, the default values are suitable in general.

(a) Video frame. (b) Low-rank component **L**. (c) Sparse component **S**.

Figure 16: Foreground/background separation of a video using `rrpca()`. Subplot (a) shows the actual video frame, which is separated into its two components. The low-rank component represents the background, and the sparse component captures the foreground objects.

performance of the `rrpca()` function using both the deterministic and the randomized SVD algorithm, which can be selected via the parameter `svdalg = c("rsvd", "svd")`. Without subspace iterations, the randomized algorithm converges after about 22 iterations. Using the deterministic SVD algorithm requires slightly fewer iterations, however, each iteration is substantially more expensive. The timings are shown in Table 3. Interestingly, the convergence profiles of the randomized algorithm appear smoother. This is due to the implicit regularization of the oversampling parameter p. Moreover, we stress that the randomized RPCA algorithm only achieves accurate results if the singular value spectrum is rapidly decaying, i.e., the data must exhibit low-rank structure.

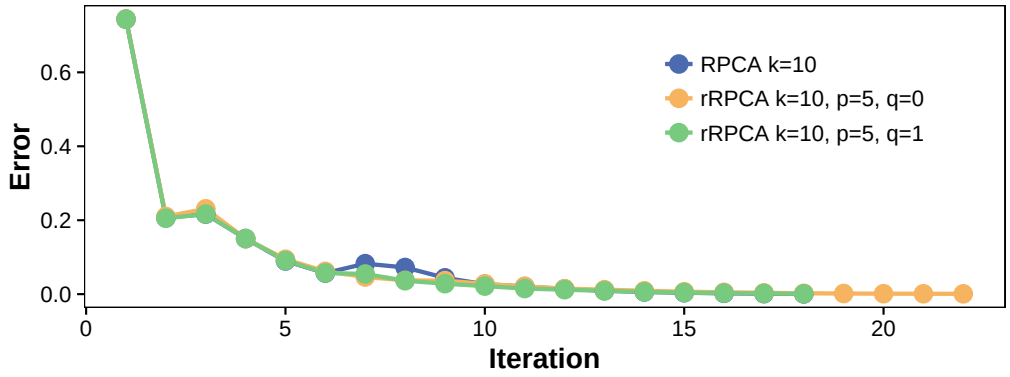

Figure 17: Convergences of the `rrpca()` function using both the deterministic and randomized SVD algorithm.

Method	Parameters	Time (s)	Speedup	Iterations	Error
`rrpca()`	`k=10, svdalg="svd"`	14.34	-	18	6.32e-08
`rrpca()`	`k=10, p=5, q=0`	6.47	2.22	22	6.85e-08
	`k=10, p=5, q=1`	7.32	1.96	19	6.52e-08

Table 3: Computational time for the deterministic and randomized RPCA algorithm.

4.5. Computational performance.

The time complexity of classic deterministic SVD algorithms is $O(mn^2)$, where it is assumed that $m \geq n$. Modern partial SVD algorithms reduce the time complexity to $O(mnk)$ (Demmel 1997). Randomized SVD, as presented here, comes also with asymptotic costs of $O(mnk)$. The key difference, however, is that the randomized SVD algorithm (wihtout subspace iterations) requires only two passes over the input matrix. Each additional subspace iteration requires one more pass over the data matrix. Hence, from a practical point of view, the algorithm is computationally more efficient.

Figure 18 shows the computational evaluation of the base `svd()`, as well as the `propack.svd()` from the **svd** package, the `svds()` from the **Rspectra** package, and the `rsvd()` function. Here two different dense random matrices are presented. In each situation the singular values decaying linearly from 1 to 0.001. The computational time and the relative reconstruction errors are computed over a sequence of different target-ranks k. In particular, the `rsvd()`

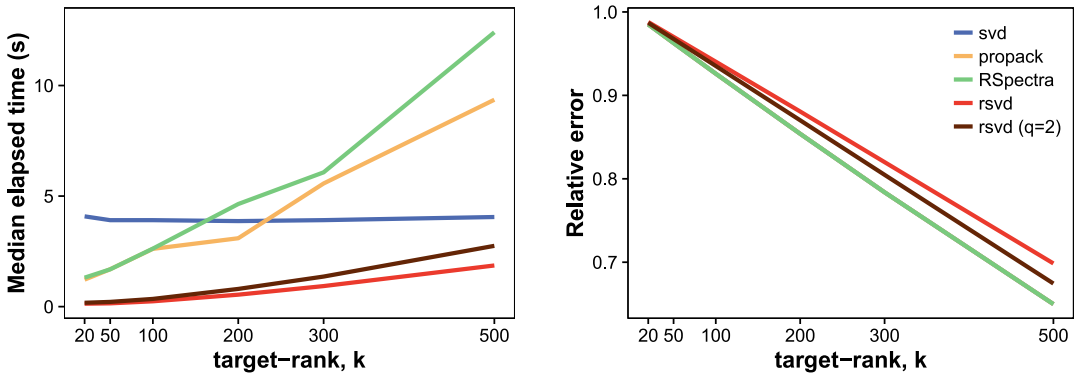

(a) Dense matrix of size 2000×2000.

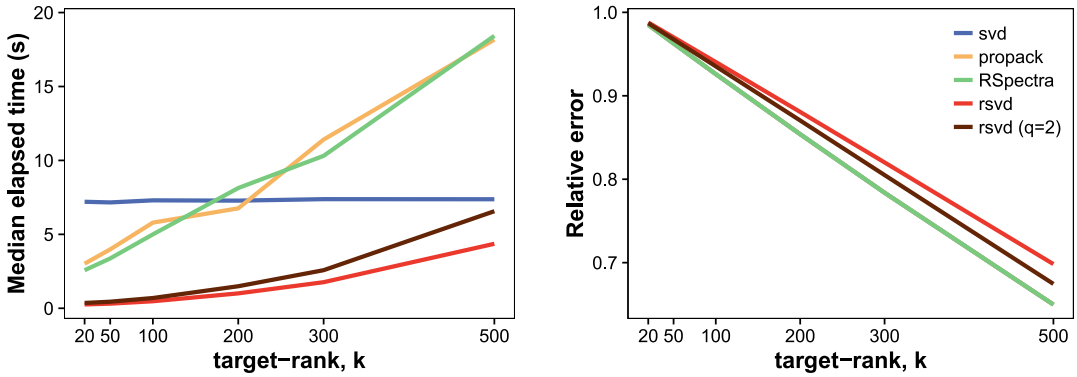

(b) Dense matrix of size 5000×2000.

Figure 18: Computational performance of singular value decomposition routines for dense matrices in R. The left column shows the median computational time and the right column shows the relative reconstruction error over varying target-ranks k. The `rsvd()` outperforms in terms of the computational time. Further, the error of the `rsvd()` algorithm can be controlled by the parameter p and q.

function achieves a substantial speedup for very low-dimensional approximations of dense matrices. The advantage of the randomized algorithm becomes pronounced with an increased matrix dimension. Hence, the randomized SVD algorithm enables the fast decomposition of large dense data matrices, while achieving competitive reconstruction errors. It is shown that that the reconstruction error can be improved by performing more subspace iterations. This allows the user to control the trade-off between computational time and accuracy.

Figure 19 shows the computational evaluation on a sparse matrix with only 5% non-zero elements. The two partial SVD algorithms are specifically designed for large sparse and structured matrices, and show here a significant computational improvement over the dense matrix of the same dimension. Although, the `rsvd()` algorithm is still competitive, the advantage of the `propack.svd()` and `svds()` functions is notable with about 1% non-zero elements for computing the dominant singular values and vectors.

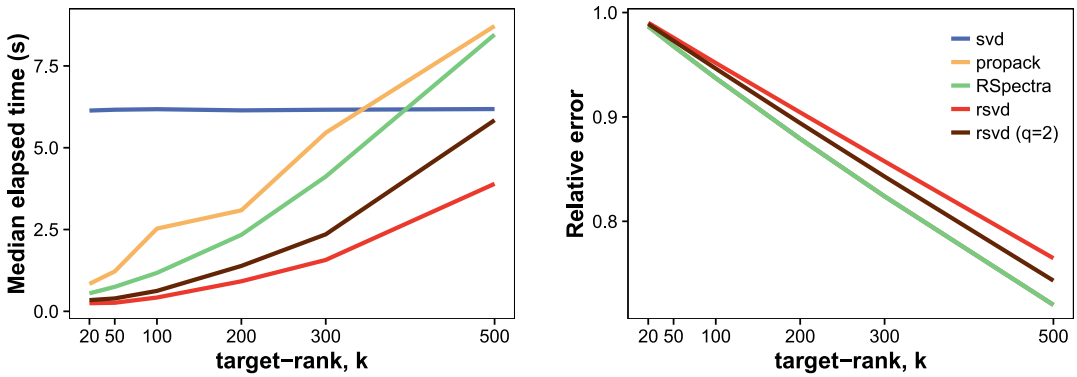

(a) Sparse matrix of size 5000 × 2000, with 5% non-zero entries.

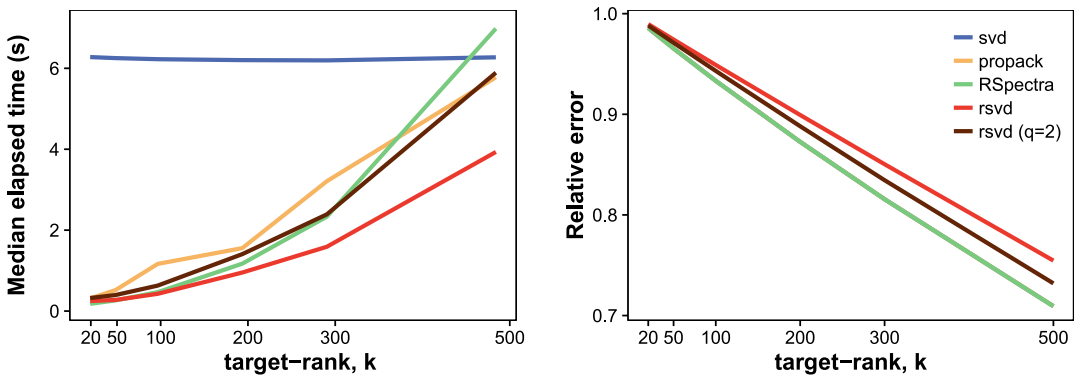

(b) Very sparse matrix of size 5000 × 2000, with 1% non-zero entries.

Figure 19: Computational performance of singular value decomposition routines for sparse matrices in R. The left column shows the median computational time and the right column shows the relative reconstruction error over varying target-ranks k. The two partial SVD algorithms `propack.svd()` and `svds()` are competitive for computing the dominant singular values and vectors of large sparse or structured matrices.

5. Conclusion

Due to the tremendous increase of high-dimensional data produced by modern sensors and social networks, data methods for dimensionality reduction are becoming increasingly important. However, despite modern computer power, massive datasets pose a tremendous computational challenge for traditional algorithms. Probabilistic algorithms can substantially ease the logistic and computational challenges in obtaining approximate matrix decompositions. This advantage becomes pronounced with an increasing matrix dimension. Randomized algorithms are feasible for even massive matrices where traditional deterministic algorithms fail. The randomized singular value decomposition is the most prominent and ubiquitous randomized algorithm. Its popularity is due to the strong theoretical error bounds and the advantage that the error can be controlled by oversampling and subspace iterations. The concept is summarized in Figure 20.

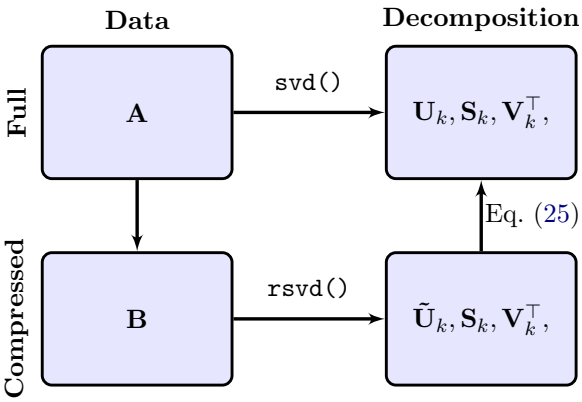

Figure 20: Conceptual architecture of the randomized singular value decomposition. The data are first compressed via right multiplication by a sampling matrix $\mathbf{\Omega}$. Next, the SVD is computed on the compressed data. Finally, the left singular vectors $\mathbf{U}_k$ may be reconstructed from the compressed singular vectors $\tilde{\mathbf{U}}_k$ by the expression in Eq. (25).

The R package **rsvd** provides computational efficient randomized routines for the singular value decomposition, principal component analysis and robust principal component analysis. The package is particularly useful for computing the approximate low rank decompositions when the rank k is substantially smaller than the matrix dimensions. The routines are intuitive to use and the performance evaluation shows that the randomized algorithm provides an efficient framework to reduce the computational demands of the traditional (deterministic) algorithms. Substantial speedups are gained over other fast (partial) SVD and PCA implements in R, while achieving a competitive reconstruction error.

The applications of the randomized matrix algorithms are ubiquitous and can be utilized for all methods relying on the computation of (generalized) eigenvalue problems. Future developments of the **rsvd** package will apply this concept to compute linear discriminant analysis, principal component regression, and canonical correlation analysis. Another important direction is to provide more efficient routines for large-scale sparse matrices using the **Matrix** package (Bates and Maechler 2016). Also additional sampling strategies based on importance sampling will be provided (Drineas and Mahoney 2016).

Acknowledgements

We would like to thank Bing Brunton for providing the formatted cropped Yale face database. JNK acknowledges support from Air Force Office of Scientific Research (FA9500-15-C-0039). SLB acknowledges support from the Department of Energy (DE-EE0006785). NBE acknowledges support from the UK Engineering and Physical Sciences Research Council (EP/L505079/1).

References

Abdi H, Williams LJ (2010). "Principal Component Analysis." *Wiley Interdisciplinary Reviews: Computational Statistics*, **2**(4), 433–459. doi:10.1002/wics.101.

Anderson E (1935). "The Irises of the Gaspe Peninsula." *Bulletin of the American Iris Society*, **59**, 2–5.

Anderson E, Bai Z, Bischof C, Blackford S, Dongarra J, Du Croz J, Greenbaum A, Hammarling S, McKenney A, Sorensen D (1999). ***LAPACK*** *Users' Guide*. 3rd edition. Society for Industrial and Applied Mathematics. doi:10.1137/1.9780898719604.

Bates D, Maechler M (2016). ***Matrix****: Sparse and Dense Matrix Classes and Methods*. R package version 1.2-6, URL https://CRAN.R-project.org/package=Matrix.

Bouwmans T, Zahzah EH (2014). "Robust PCA via Principal Component Pursuit: A Review for a Comparative Evaluation in Video Surveillance." *Computer Vision and Image Understanding*, **122**, 22–34. doi:10.1016/j.cviu.2013.11.009.

Candès EJ, Li X, Ma Y, Wright J (2011). "Robust Principal Component Analysis?" *Journal of the ACM (JACM)*, **58**(3), 11.

Cunningham JP, Ghahramani Z (2015). "Linear Dimensionality Reduction: Survey, Insights, and Generalizations." *Journal of Machine Learning Research*, **16**, 2859–2900.

Demmel J (1997). *Applied Numerical Linear Algebra*. Society for Industrial and Applied Mathematics. ISBN 9780898713893. doi:10.1137/1.9781611971446.

Donoho DL (2000). "High-Dimensional Data Analysis: The Curses and Blessings of Dimensionality." *AMS Math Challenges Lecture*, pp. 1–32.

Drineas P, Mahoney MW (2016). "RandNLA: Randomized Numerical Linear Algebra." *Communications of the ACM*, **59**(6), 80–90. doi:10.1145/2842602.

Eckart C, Young G (1936). "The Approximation of one Matrix by Another of Lower Rank." *Psychometrika*, **1**(3), 211–218.

Erichson NB, Donovan C (2016). "Randomized Low-Rank Dynamic Mode Decomposition for Motion Detection." *Computer Vision and Image Understanding*, **46**, 40–50. doi:10.1016/j.cviu.2016.02.005.

Fienup JR (1997). "Invariant Error Metrics for Image Reconstruction." *Applied Optics*, **36**(32), 8352–8357.

Frieze A, Kannan R, Vempala S (2004). "Fast Monte-Carlo Algorithms for Finding Low-Rank Approximations." *Journal of the ACM*, **51**(6), 1025–1041.

Gavish M, Donoho DL (2014). "The Optimal Hard Threshold for Singular Values is." *Information Theory, IEEE Transactions on*, **60**(8), 5040–5053.

Georghiades AS, Belhumeur PN, Kriegman DJ (2001). "From Few to Many: Illumination Cone Models for Face Recognition Under Variable Lighting and Pose." *IEEE transactions on pattern analysis and machine intelligence*, **23**(6), 643–660.

Golub G, Kahan W (1965). "Calculating the Singular Values and Pseudo-Inverse of a Matrix." *Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis*, **2**(2), 205–224.

Golub GH, Reinsch C (1970). "Singular Value Decomposition and Least Squares Solutions." *Numerische Mathematik*, **14**(5), 403–420.

Golub GH, Van Loan CF (1996). *Matrix Computations*. 3 edition. Johns Hopkins University Press, Baltimore, MD, USA. ISBN 0801854148.

Goyette N, Jodoin PM, Porikli F, Konrad J, Ishwar P (2012). "Changedetection.net: A new Change Detection Benchmark Dataset." In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2012 IEEE Computer Society Conference on*, pp. 1–8. IEEE.

Greenacre MJ (2010). *Biplots in Practice*. Fundacion BBVA.

Gu M (2015). "Subspace Iteration Randomization and Singular Value Problems." *SIAM Journal on Scientific Computing*, **37**(3), 1139–1173.

Halko N, Martinsson PG, Shkolnisky Y, Tygert M (2011a). "An Algorithm for the Principal Component Analysis of Large Data Sets." *SIAM Journal on Scientific Computing*, **33**(5), 2580–2594.

Halko N, Martinsson PG, Tropp JA (2011b). "Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions." *SIAM Review*, **53**(2), 217–288. doi:10.1137/090771806.

Hastie T, Tibshirani R, Friedman J (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics, 2nd edition. Springer-Verlag. ISBN 9780387216065.

Izenman AJ (2008). *Modern Multivariate Statistical Techniques: Regression, Classification, and Manifold Learning*. Springer-Verlag. ISBN 0387781889.

Jolliffe I (2002). *Principal Component Analysis*. Springer Series in Statistics, 2nd edition. Springer-Verlag. ISBN 9780387954424.

Kirby M, Sirovich L (1990). "Application of the Karhunen-Loeve Procedure for the Characterization of Human Faces." *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, **12**(1), 103–108.

Korobeynikov A, Larsen RM (2016). ***svd**: Interfaces to Various State-of-Art SVD and Eigensolvers*. R package version 0.4, URL https://CRAN.R-project.org/package=svd.

Larsen RM (1998). "Lanczos Bidiagonalization with Partial Reorthogonalization." *DAIMI Report Series*, **27**(537), 1–101. doi:10.7146/dpb.v27i537.7070.

Lehoucq R, Sorensen D, Yang C (1998). ***ARPACK** Users' Guide*. Society for Industrial and Applied Mathematics. doi:10.1137/1.9780898719628.

Lin Z, Chen M, Ma Y (2010). "The Augmented Lagrange Multiplier Method for Exact Recovery of Corrupted Low-Rank Matrices." *arXiv preprint*, pp. 1–23. http://arxiv.org/abs/1009.5055.

Mahoney MW (2011). "Randomized Algorithms for Matrices and Data." *Foundations and Trends in Machine Learning*, **3**(2), 123–224. doi:10.1561/2200000035.

Martinsson PG (2016). "Randomized Methods for Matrix Computations and Analysis of High Dimensional Data." *arXiv preprint*, pp. 1–55. https://arxiv.org/abs/1607.01649.

Martinsson PG, Rokhlin V, Tygert M (2011). "A Randomized Algorithm for the Decomposition of Matrices." *Applied and Computational Harmonic Analysis*, **30**(1), 47–68. doi:10.1016/j.acha.2010.02.003.

Martinsson PG, Voronin S (2015). "A randomized Blocked Algorithm for Efficiently Computing Rank-Revealing Factorizations of Matrices." *arXiv preprint*, pp. 1–26. http://arxiv.org/abs/1503.07157.

Mersmann O, Beleites C, Hurling R, Friedman A (2015). ***microbenchmark**: Accurate Timing Functions*. R package version 1.4-2.1, URL http://CRAN.R-project.org/package=microbenchmark.

Murphy KP (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press. ISBN 9780262018029.

Pearson K (1901). "On Lines and Planes of Closest Fit to Systems of Points in Space." *Philosophical Magazine Series 6*, **2**(11), 559–572. doi:10.1080/14786440109462720.

Qiu Y, Mei J, Guennebaud G, Niesen J (2016). ***RSpectra**: Solvers for Large Scale Eigenvalue and SVD Problems*. R package version 0.12-0, URL https://CRAN.R-project.org/package=RSpectra.

Rivasplata O (2012). "Subgaussian Random Variables: An Expository Note." *Internet publication, PDF*. URL http://www.stat.cmu.edu/~arinaldo/36788/subgaussians.pdf.

Rokhlin V, Szlam A, Tygert M (2009). "A Randomized Algorithm for Principal Component Analysis." *SIAM Journal on Matrix Analysis and Applications*, **31**(3), 1100–1124.

Sarlos T (2006). "Improved Approximation Algorithms for Large Matrices via Random Projections." In *Foundations of Computer Science. 47th Annual IEEE Symposium on*, pp. 143–152.

Stewart GW (1993). "On the Early History of the Singular Value Decomposition." *SIAM review*, **35**(4), 551–566.

Szlam A, Kluger Y, Tygert M (2014). "An Implementation of a Randomized Algorithm for Principal Component Analysis." *arXiv preprint arXiv:1412.3510*.

Turk MA, Pentland AP (1991). "Face Recognition using Eigenfaces." In *Proceedings on Computer Vision and Pattern Recognition*, pp. 586–591. IEEE.

Venables WN, Ripley BD (2002). *Modern Applied Statistics with S-PLUS*. Statistics and Computing. Springer-Verlag. ISBN 9780387954578. doi:10.1007/978-0-387-21706-2.

Voronin S, Martinsson PG (2015). "**RSVDPACK**: Subroutines for Computing Partial Singular Value Decompositions via Randomized Sampling on Single Core, Multi Core, and GPU Architectures." *arXiv preprint*, pp. 1–15. http://arxiv.org/abs/1502.05366.

Watkins DS (2002). *Fundamentals of Matrix Computations*. 2 edition. John Wiley & Sons. ISBN 9780470528334. doi:10.1002/0471249718.

Wickham H (2009). ***ggplot2**: Elegant Graphics for Data Analysis*. Springer-Verlag. ISBN 9780387981406. URL http://had.co.nz/ggplot2/book.

Woolfe F, Liberty E, Rokhlin V, Tygert M (2008). "A Fast Randomized Algorithm for the Approximation of Matrices." *Applied and Computational Harmonic Analysis*, **25**(3), 335–366.

Wright J, Ganesh A, Rao S, Peng Y, Ma Y (2009). "Robust Principal Component Analysis: Exact Recovery of Corrupted Low-Rank Matrices via Convex Optimization." In *Advances in Neural Information Processing Systems*, pp. 2080–2088.

Affiliation:

N. Benjamin Erichson
School of Mathematics and Statistics, University of St Andrews
St Andrews, UK
E-mail: nbe@st-andrews.ac.uk

Sergey Voronin
Department of Mathematics, Tufts University
Medford, Massachusetts, U.S.

Steven L. Brunton
Department of Mechanical Engineering, University of Washington
Seattle, Washington, U.S.

J. Nathan Kutz
Department of Applied Mathematics, University of Washington
Seattle, Washington, U.S.